Some thoughts on education and political priorities

Summary

Although we understand some systems well enough to make precise or statistical predictions, most interesting systems - whether physical, mental, cultural, or virtual - are complex, nonlinear, and have properties that emerge from feedback between many interactions. Exhaustive searches of all possibilities are impossible. Unfathomable and unintended consequences dominate. Problems cascade. Complex systems are hard to understand, predict and control.

A growing fraction of the world has made a partial transition from a) small, relatively simple, hierarchical, primitive, zero-sum hunter-gatherer tribes based on superstition (almost total ignorance of complex systems), shared aims, personal exchange and widespread violence, to b) large, relatively complex, decentralised, technological, nonzero-sum market-based cultures based on science (increasingly accurate predictions and control in some fields), diverse aims, impersonal exchange, trade, private property, and (roughly) equal protection under the law. Humans have made transitions from numerology to mathematics, from astrology to astronomy, from alchemy to chemistry, from witchcraft to neuroscience, from tallies to quantum computation. However, while our ancestor chiefs understood bows, horses, and agriculture, our contemporary chiefs (and those in the media responsible for scrutiny of decisions) generally do not understand their equivalents, and are often less experienced in managing complex organisations than their predecessors.

The education of the majority even in rich countries is between awful and mediocre. In England, few are well-trained in the basics of extended writing or mathematical and scientific modelling and problem-solving. Less than 10 percent per year leave school with formal training in basics such as exponential functions, 'normal distributions' ('the bell curve'), and conditional probability. Less than one percent are well educated in the basics of how the 'unreasonable effectiveness of mathematics' provides the language of nature and a foundation for our scientific civilisation. Only a small subset of that <1% then study trans-disciplinary issues concerning complex systems. This number has approximately zero overlap with powerful decision-makers. Generally, they are badly (or narrowly) educated and trained (even elite universities offer courses that are thought to prepare future political decision-makers but are clearly inadequate and in some ways damaging). They also usually operate in institutions that have vastly more ambitious formal goals than the dysfunctional management could possibly achieve, and which generally select *for* the worst aspects of chimp politics and *against* those skills seen in rare successful organisations (e.g the ability to *simplify, focus*, and *admit errors*). Most politicians, officials, and advisers operate with fragments of philosophy, little knowledge of maths or science ([few MPs can answer](#) even simple probability questions yet most are confident in their judgement), and little experience in well-managed complex organisations. The skills, and approach to problems, of our best mathematicians, scientists, and entrepreneurs are almost totally shut out of vital decisions. We do not have a problem with 'too much cynicism' - we have a problem with *too much trust in people and institutions that are not fit to control so much.*

The consequences are increasingly dangerous as markets, science and technology disrupt all existing institutions and traditions, and enhance the dangerous potential of our evolved nature to inflict huge physical destruction and to manipulate the feelings and ideas of many people (including, sometimes *particularly*, the best educated) through 'information operations'. Our fragile civilisation is vulnerable to large shocks and a continuation of traditional human politics as it was during 6 million years of hominid evolution – an attempt to secure in-group cohesion, prosperity and strength in order to dominate or destroy nearby out-groups in competition for scarce resources - could kill billions. We need big changes to schools, universities, and political and other institutions

for their own sake and to help us limit harm done by those who, entangled with trends described below, pursue dreams of military glory, 'that attractive rainbow that rises in showers of blood.'

Some ideas are presented, aimed mainly at 15-25 year-olds, for what physicist Murray Gell Mann described as an 'Odyssean' education synthesising a) maths and the natural sciences, b) the social sciences, and c) the humanities and arts, into crude, trans-disciplinary, integrative thinking. This should combine courses like *The Big History Project*, Berkeley's '*Physics for Future Presidents*' (or Professor Timothy Gowers' planned equivalent for maths) with the best of the humanities; add new skills such as coding; and give training in managing complex projects and using modern tools (e.g agent-based models, ABMs). Universities should develop alternatives to *Politics, Philosophy, and Economics* such as *Ancient and Modern History, Physics for Future Presidents, and Programming*. We need leaders with an understanding of Thucydides and statistical modelling, who have read *The Brothers Karamazov* and *The Quark and the Jaguar*, who can feel Kipling's Kim and succeed in Tetlock's *Good Judgement Project*. An Odyssean education would focus on humans' biggest and most important problems and explain connections between them to train *synthesisers*. An approach is suggested here based on seven broad fields with some big, broad goals.

1. Maths and complexity. Solve the Millennium Problems, better prediction of complex networks.
2. Energy and space. Ubiquitous cheap energy and opening space for science and commerce.
3. Physics and computation. Exploration beyond the Standard Model of particle physics, better materials and computers, digital fabrication, and quantum computation.
4. Biological engineering. Understanding the biological basis of personality and cognition, personalised medicine, and computational and synthetic biology.
5. Mind and machine. Quantitative models of the mind and machine intelligence applications.
6. The scientific method, education, training and decisions. Nielsen's vision of decentralised coordination of expertise and data-driven intelligence ('a scientific social web that directs scientists' attention where it is most valuable'); more ambitious and scientifically tested personalised education; training and tools that measurably improve decisions (e.g. ABMs).
7. Political economy, philosophy, and avoiding catastrophes. Replacements for failed economic ideas and traditional political philosophies; new institutions (e.g. new civil service systems and international institutions, a UK DARPA and TALPIOT (non-military), decentralised health services).

Such an education and training might develop synthesisers who have 1) a crude but useful grasp of connections between the biggest challenges based on trans-disciplinary thinking about complex systems; 2) a cool Thucydidean courage to face reality including their own errors and motives; 3) the ability to take better decisions and *adapt* fast to failures; 4) an evolutionary perspective on complex systems and institutional design (rather than the typical Cartesian 'chief of the tribe' perspective); and 5) an ability to shape new institutions operating like an immune system that will bring better chances to avoid, survive, and limit damage done by inevitable disasters.

Focus is hard to hold in politics. After 1945, Dean Acheson quipped that Britain had failed to find a post-imperial role. It is suggested here that *this role should focus on making ourselves the leading country for education and science*: Pericles described Athens as 'the school of Greece', we could be the school of the world. Who knows what would happen to a political culture if a party embraced education and science as its defining mission and *therefore changed the nature of the people running it and the way they make decisions and priorities*. We already have a head start; we lack focus. Large improvements in education and training are easier to achieve than solving many other big problems and will contribute to their solution. Progress *could* encourage non-zerosum institutions and global cooperation - alternatives to traditional politics and destruction of competitors. However, the spread of knowledge and education is itself a danger and cannot eliminate gaps in wealth and power created partly by unequally distributed heritable characteristics.

Earthrise from the moon, 1968

'What we are creating now is a monster whose influence is going to change history, provided there is any history left. Yet it would be impossible not to see it through... The world could be conquered, but this nation of puritans will not grab its chance; we will be able to go into space way beyond the moon if only people could keep pace with what they create.' Von Neumann on nuclear and computer technology, 1945.

'I feel it myself, the glitter of nuclear weapons. It is irresistible if you come to them as a scientist. To feel it's there in your hands. To release the energy that fuels the stars. To let it do your bidding. And to perform these miracles, to lift a million tons of rock into the sky, it is something that gives people an illusion of illimitable power, and it is in some ways responsible for all our troubles, I would say, this is what you might call 'technical arrogance' that overcomes people when they see what they can do with their minds.' Freeman Dyson.

'We've arranged a global civilization in which most crucial elements - transportation, communications, and all other industries; agriculture, medicine, education, entertainment, protecting the environment; and even the key democratic institution of voting - profoundly depend on science and technology. We have also arranged things so that almost no one understands science and technology. This is a prescription for disaster. We might get away with it for a while, but sooner or later this combustible mixture of ignorance and power is going to blow up in our faces.' Carl Sagan.

'We are all in a situation that resembles driving a fast vehicle at night over unknown terrain that is rough, full of gullies, with precipices not far off. Some kind of headlight, even a feeble and flickering one, may help to avoid some of the worst disasters.' John Holland, inventor of 'genetic algorithms'.

Introduction: Integrative thinking - 'a crude look at the whole'

'The love of complexity without reductionism makes art; the love of complexity with reductionism makes science.' E.O. Wilson.

'Information consumes ... the attention of its recipients. Hence a wealth of information creates a poverty of attention and a need to allocate that attention efficiently among the overabundance of information sources that might consume it... Attention is generally scarce in organizations. It is particularly scarce at the top of organizations. And it is desperately scarce at the top of the organization we call the United States Government. There is only one President and ... a frightening array of matters converges on that single, serial information processing system... In a knowledge-rich world, progress does not lie in the direction of reading information faster, writing it faster, and storing more of it. Progress lies in the direction of extracting and exploiting the patterns of the world... And that progress will depend on ... our ability to devise better and more powerful thinking programs for man and machine.' Herbert Simon (*Designing Organizations for an Information-rich World*, 1969).[1]

This essay is aimed mainly at ~15-25 year-olds and those interested in more ambitious education and training for them. Not only are most of them forced into mediocre education but they are also then forced into dysfunctional institutions where many face awful choices: either conform to the patterns set by middle-aged mediocrities (don't pursue excellence, don't challenge bosses' errors, and so on) or soon be despised and unemployed. Some of the ideas sketched here may help some to shape their own education and create their own institutions. As people such as Linus Torvald (Linux) or Mark Zuckerberg (Facebook) have shown, the young are capable of much more than the powerful and middle-aged, who control so many institutions, like to admit.[2] A significant change in education and training could also help us partly ameliorate the grim cycle of predictable political dysfunction.

Although it is generally psychologically unpleasant to focus on our problems and admit the great weaknesses of our institutions, including what we in politics and our leaders do not understand, it is only through honest analysis that progress is possible. As Maxwell said, 'Thoroughly conscious ignorance is a prelude to every real advance in knowledge.'

Reliable knowledge about what works (how do people learn, what teaching methods work, how to use technology) should be built cumulatively on the foundation of empirical tests as suggested by physicist Carl Wieman and others (cf. Section 6 and Endnote). New systems and curricula would work in different ways for a school, a University, or an equivalent to a boxing gym for politicians if such a thing existed.

In particular, it is suggested that we need *an 'Odyssean' education* so that a substantial fraction of teenagers, students and adults might understand something of our biggest intellectual and practical problems, and be trained to take effective action.

---

[1] In the final sentence Simon proposed a goal of 'an order of magnitude increase by 1980 in the speed with which a human being can learn a difficult school subject - say, a foreign language or arithmetic.'

[2] I have seen many organisations in which 'junior' people aged 20-30, paid a fraction of the boss, routinely do more to solve the organisation's problems (and suppress failure) in an hour of disciplined action per day than do bad senior executives, paid £100-300,000, over days and months. Middle-aged middle management try hard to prevent such people rising, hence they leave, hence the overall culture deteriorates as the people responsible for failure rise with promotions and pay rises.

What does an 'Odyssean' education mean?

The Nobel-winning physicist, [Murray Gell Mann](), one of the architects of the Standard Model of particle physics and namer of the 'quark',[3] has described a scientific and political need for an 'Odyssean' philosophy that can synthesise a) maths and the natural sciences, b) the social sciences, and c) the humanities and arts, into necessarily crude, trans-disciplinary, integrative thinking about complex systems.

> 'Today the network of relationships linking the human race to itself and to the rest of the biosphere is so complex that all aspects affect all others to an extraordinary degree. Someone should be studying the whole system, however crudely that has to be done, because no gluing together of partial studies of a complex nonlinear system can give a good idea of the behavior of the whole...

> 'Those who study complex adaptive systems are beginning to find some general principles that underlie all such systems, and seeking out those principles requires intensive discussions and collaborations among specialists in a great many fields. Of course the careful and inspired study of each specialty remains as vital as ever. But integration of those specialities is urgently needed as well. Important contributions are made by the handful of scholars and scientists who are transforming themselves from specialists into students of simplicity and complexity or of complex adaptive systems in general...

> '[There is] the distinction (made famous by Nietzsche) between "Apollonians", who favor logic, the analytical approach, and a dispassionate weighing of the evidence, and "Dionysians'", who lean more toward intuition, synthesis, and passion…[4] But some of us seem to belong to another category: the "Odysseans", who combine the two predilections in their quest for connections among ideas… We need to celebrate the contribution of those who dare take what I call "a crude look at the whole"…

> '… broadly integrative thinking is relegated to cocktail parties. In academic life, in bureaucracies, and elsewhere, the task of integration is insufficiently respected. Yet anyone at the top of an organization … has to make decisions *as if* all aspects of a situation, along with the interaction among those aspects, were being taken into account. Is it reasonable for the leader, reaching down into the organization for help, to encounter specialists and for integrative thinking to take place only when he or she makes the final intuitive judgements?

> '[A] multiplicity of crude but integrative policy studies, involving not just linear projection but evolution and highly nonlinear simulation and gaming, may provide some modest help in generating a collective foresight function for the human race…

> 'Given the immense complexity of the numerous interlocking issues facing humanity, foresight demands the ability to identify and gather great quantities of relevant information; the ability to catch glimpses, using that information, of the choices offered by the branching alternative histories of the future, and the wisdom to select simplifications and approximations that do not sacrifice the representation of critical qualitative issues, especially issues of values…

---

[3] Gell Mann got the name 'quark' from a line in *Finnegans Wake*, 'Three quarks for Muster Mark'.

[4] In Nietzsche's metaphor, there is the realm of Apollo – order, reason, sculpture – and the world of Dionysus – chaos, instinct, music; man is 'creature and creator', 'Homer *against* Plato, that is the genuine antagonism.' Aeschylus' *Oresteia* shows Apollo dominant; Euripides' *Baccae* shows Dionysus dominant.

'Computers ... can serve us both by learning or adapting themselves and by modelling or simulating systems in the real world that learn or adapt or evolve. .. Powerful computers are essential for assistance in looking into the future, but we must not allow their use to bias the formulation of problems toward the quantifiable and analyzable at the expense of the important.'[5]

The biologist E.O. Wilson, one of the founders of modern evolutionary biology, has also described the need to integrate knowledge across subjects. He describes 'the Ionian Enchantment' - the belief in the unity of nature and the search for the principles that explain it. 'Its central tenet is the unification of knowledge', it has been the foundation for modern science, and it has been consistently vindicated. Wilson argues that what he calls 'consilience' provides an 'Ariadne's Thread' of explanation from biochemisty to genetics to quantitative models of the brain to culture.

> 'Consilience is the key to unification... [It is] a "jumping together" of knowledge by the linking of facts and fact-based theory across disciplines to create a common groundwork of explanation… [It] gives ultimate purpose to intellect. It promises that order, not chaos, lies beyond the horizon...
>
> 'Given that human action comprises events of physical causation, why should the social sciences and humanities be impervious to consilience with the natural sciences? ... Nothing fundamental separates the course of human history from the course of physical history, whether in the stars or in organic diversity. Astronomy, geology, and evolutionary biology are examples of primarily historical disciplines linked by consilience to the rest of the natural sciences. [I]f ten thousand humanoid histories could be traced to ten thousand Earthlike planets, and from comparative study of those histories empirical tests and principles evolved, historiography – the explanation of historical trends – would already be a natural science.'

Wilson argues that Consilience will revive the original hopes of the Enlightenment that science would explain human culture - hopes that crumbled because of science's inability to explain the mind and the emotional aspects of humans that Romanticism explored. The enterprises of culture will, he thinks, gradually separate into (i) maths and the natural sciences and (ii) the humanities (particularly the creative arts). The social sciences will divide in each direction. Philosophy, the study of what we do not know scientifically and the reasons why, will gradually shrink. The humanities and natural sciences will gradually become closer and partly fuse. Even if one does not agree with this broad idea,[6] and even though it is clear that much progress must and will continue to be made by developing traditional specialist subjects (trans-disciplinary fads can themselves be dangerous), one can agree with Gell Mann about the urgent need for 'integrators' with the education to take 'a crude look at the whole' and with Wilson about the need for 'synthesizers'.

> 'We are drowning in information, while starving for wisdom. The world henceforth will be run by synthesizers, people able to put together the right information at the right time, think critically about it, and make important choices wisely.' (E.O. Wilson)

One of the things that 'synthesizers' need to learn is the way that many themes cut across subjects and generate new subjects. Trans-disciplinary studies of complex systems have been profoundly affected by connected intellectual revolutions in physics, maths, logic, computation, and biology, though these connections have so far barely been integrated in school or university curricula. Ideas

---

[5] Cf. Gell Mann's *The Quark and the Jaguar*.

[6] 'The search for consilience might seem at first to imprison creativity. The opposite is true. A united system of knowledge is the surest means of identifying the still unexplored domains of reality. It provides a clear map of what is known, and it frames the most productive questions for future inquiry.' E.O.W

about 'information' cut across physics (thermodynamics and entropy),[7] computation (bits, qubits, and 'artificial agents'), economics ('economic agents', Hayek's description of prices as an 'information discovery process'), evolution (genetic networks),[8] the brain (neuronal networks), 'intelligence failures'[9] and other subjects. Physics is inseparable from old philosophical questions (it is 'experimental metaphysics') and it is merging with computer science to produce quantum computation and 'quantum information theory'.[10] Evolutionary ideas took hold in economics (Hume, Smith) and biology (Darwin), and they have now been incorporated into computer science (e.g. 'genetic algorithms') in order to 'evolve' solutions to problems in large search spaces, and they have suggested ideas for engineering solutions. (For example, it is suggested that dangers such as bioterrorism or pandemics should be defended by developing 'artificial immune systems' in which defences operate according to the evolutionary principles of i) generating lots of solutions with random variation, and ii) differential replication of the most effective agents, instead of reliance on traditional centralised institutions that make similar mistakes repeatedly.) Machine intelligence (or 'automated reasoning') has been shaped by, and in turn is reshaping, ideas about the mind and is now used to design computers, including those that are used to investigate the mind. Behavioural genetics, evolutionary psychology, cognitive science and neuroscience are reshaping not only economics (e.g. fMRI scans of people playing financial games) and history (e.g. the influence of evolved antipathy for out-groups) but also how we design institutions (e.g. how we evolved to succeed in the sexual politics of small, primitive collectivist tribes hence many of the standard features of internal politics). As will be discussed, there are various developments in education that seek to reflect these trans-disciplinary themes: e.g. the Nobel-winning neuroscientist, Eric Kandel, plans a new PhD programme combining neuroscience, psychology and art history as part of the 'Mind, Brain, Behaviour Project' at Columbia.[11]

An Odyssean curriculum would give students and politicians some mathematical foundations and a map to navigate such subjects without requiring a deep specialist understanding of each element. In order to provide some structure to such an enterprise, *a schema of seven big areas* and some material is sketched (links, Endnotes from page 129, Reading List).

---

[7] Ideas linking physics and information began in the 19th Century with thermodynamics and the infamous 'Maxwell's Demon'. Shannon's '*A Mathematical Theory of Communication*' (1948) is regarded as the founding text of information theory and introduced the idea of treating information like entropy in physics as a measure of disorder or uncertainty. 'My greatest concern was what to call it. I thought of calling it 'information', but the word was overly used, so I decided to call it 'uncertainty'. When I discussed it with John von Neumann, he had a better idea. Von Neumann told me, "You should call it entropy, for two reasons. In the first place your uncertainty function has been used in statistical mechanics under that name, so it already has a name. In the second place, and more important, nobody knows what entropy really is, so in a debate you will always have the advantage."' Crutchfield uses Shannon's ideas of information theory to analyse the scientific process of extracting useful information about physical systems.

[8] The human genome consists of ~3.2 billion 'base pairs' (or 'bases') written in a language of four DNA nucleotides (ACGT). Each 'base pair' means one of the four letters ACGT on one strand of the double helix plus its partner on the other strand. A section of pairs is a 'gene' (the definition of 'gene' remains controversial). Since a bit (binary digit) halves prior uncertainty (i.e. a bit answers one yes/no question), a bit encodes two possibilities. Each base holds 2 bits so the human genome has an 'information content' of ~6.4 gigabits, which ≈ 800 megabytes of information (uncompressed).

[9] Roberta Wohlstetter's classic '*Pearl Harbor: Warning and Decision*' applied Shannon's distinction between 'signal' and 'noise' in information theory to cases of intelligence failure.

[10] At the intersection of information theory, quantum mechanics and game theory is 'quantum game theory'. At the intersection of biology and quantum mechanics is 'quantum biology'; e.g. a recent paper concludes that the evidence concerning photosynthesis and magnetoreception 'is in favor of both these systems not only containing quantum coherence in a hot and wet biological environment, but also that it is used to gain a biological advantage.'

[11] Cf. Kandel op-ed in NYT and '*The artist and the machine*', Michael Nielsen (2013).

E.O. Wilson has written about the importance of education focusing on the biggest questions.

> 'But another way of teaching science, which I adopted during forty years of teaching at Harvard, is to instruct from the top down... By that I mean that you don't start with elements like calculus or analytic geometry… You start from the top ... with the big topics that mean something immediate and important to people. For example: What is life? What's the meaning of life? What is the meaning of sex? Why do we have to die? What is the meaning of the aging process? And so on. When you've got the attention of the audience, then you break the big questions down into subsidiary questions that involve scientific subjects.

> 'I found at Harvard that I could take mathophobes and students who had very little interest in science, and by examining a subject like the meaning of sex and breaking it down I could soon have the whole class deriving a basic equation from first principles in population genetics and pondering at considerable length the chemical basis of the genetic code. If I'd started the other way, from the bottom up, I think I would have lost half the class in the first couple of weeks. So that's one way to increase interest in science: make it immediate, personal, and interesting by proceeding from the top down with questions that students really care about and understand intuitively from the start.'

This is the approach I have taken. I loved studying classics, and do not think there is a better book to study than Thucydides as training for politics, but I have focused on maths and sciences because, first, there is far more understanding about how to teach humanities successfully than maths, science, and crude interdisciplinary integrative thinking about big issues; second, because politics is dominated by people with little knowledge of many of these subjects. This paper is very crude and written by someone with no relevant expertise in any subject except politics - it is a bad version of something I wish I had been given aged 15, prompted partly by thinking, after reading about the Russian Kolmogorov schools, 'we need some schools like that... what would the pupils study?' The real need is a detailed online map for such an education that is gradually filled in by people with proper understanding of the individual topics and of teaching.[12]

*

---

[12] It would be more beneficial to follow the links and Reading List than to study this author's argument. There is little original here and the original bits are the least valuable. I am circulating it to get feedback before distributing wider. Please send thoughts, corrections etc to theodysseanproject@gmail.com, and/or follow @odysseanproject.

1. Maths, complexity, and prediction. *Can we make breakthroughs in understanding, prediction and control of complex nonlinear systems?*

'All great things have their origin in that which is small.' Lao Tzu.

> *'What can be avoided*
> *Whose end is purposed by the mighty gods?*
> *Yet Caesar shall go forth, for these predictions*
> *Are to the world in general as to Caesar.'*
> *Julius Caesar, II.2.*

'Ideas thus made up of several simple ones put together, I call Complex; such as are Beauty, Gratitude, a Man, an Army, the Universe.' Locke.

'I can calculate the motion of heavenly bodies but not the madness of people.' Newton, after the South Sea Bubble 'Ponzi scheme'.

'Everything in war is very simple, but the simplest thing is difficult. The difficulties accumulate and end by producing a kind of friction that is inconceivable unless one has experienced war… Countless minor incidents - the kind you can never really foresee - combine to lower the general level of performance, so that one always falls short of the intended goal. Iron will-power can overcome this friction … but of course it wears down the machine as well… Friction is the only concept that … corresponds to the factors that distinguish real war from war on paper. The … army and everything else related to it is basically very simple and therefore seems easy to manage. But … each part is composed of individuals, every one of whom retains his potential of friction… This tremendous friction … is everywhere in contact with chance, and brings about effects that cannot be measured… Friction … is the force that makes the apparently easy so difficult… Finally … all action takes place … in a kind of twilight, which like fog or moonlight, often tends to make things seem grotesque and larger than they really are. Whatever is hidden from full view in this feeble light has to be guessed at by talent, or simply left to chance.' Clausewitz.

'It is a wonderful feeling to recognise the unity of complex phenomena that to direct observation appear to be quite separate things.' Einstein to Grossman, 1901.

'All stable processes we shall predict. All unstable processes we shall control.' Von Neumann.

'Imagine how much harder physics would be if electrons had feelings.' Richard Feynman.

At the beginning of *From Russia With Love* (the movie not the book), Kronsteen, a Russian chess master and SPECTRE strategist, is summoned to Blofeld's lair to discuss the plot to steal the super-secret 'Lektor Decoder' and kill Bond. Kronsteen outlines to Blofeld his plan to trick Bond into stealing the machine for SPECTRE.

> Blofeld: Kronsteen, you are sure this plan is foolproof?
> Kronsteen: Yes it is, because I have anticipated every possible variation of counter-move.

Political analysis is full of chess metaphors, reflecting an old tradition of seeing games as models of physical and social reality.[13] A game which has ten different possible moves at each turn and runs for two turns has 10^2 possible ways of being played; if it runs for fifty turns it has 10^{50} possible ways of being played, 'a number which substantially exceeds the number of atoms in the whole of our planet earth' (Holland); if it runs for ninety turns it has 10^{90} possible ways of being played,

---

[13] 'Time is a child moving counters in a game; the royal power is a child's' (Heraclitus).

which is about the estimated number of atoms in the Universe. Chess is merely 32 pieces on an 8x8 grid with a few simple rules but the number of possible games is much greater than 10^{90}.

Many practical problems (e.g logistics, designing new drugs) are equivalent to the Travelling Salesman Problem. For any Travelling Salesman Problem (TSP) involving travelling to n cities, the number of possible tours when starting with a specific city is: (n-1)!/2.[14] For 33 cities, the total number of possible journeys is:

$$32!/2 = 131,565,418,466,846,765,083,609,006,080,000,000$$

The IBM Roadrunner, the faster supercomputer in the world in 2009, could perform 1,457 trillion operations per second. If we could arrange the tours such that examining each one would take only one arithmetical operation, then it would take it ~28 trillion years to examine all possible routes between 33 cities, about twice the estimated age of the Universe. As *n* grows *linearly* (add one city, add another etc), the number of possible routes grows *exponentially*. The way in which the number of possible options scales up *exponentially* as the number of agents scales up *linearly*, and the difficulty of finding solutions quickly in vast search landscapes, connects to one of the most important questions in maths and computer science, the famous $1 million dollar 'P=NP?' Clay Millennium Prize (see Endnote on Computation).

Kronsteen's confidence, often seen in politics, is therefore misplaced even in chess. It is far beyond our ability to anticipate 'every possible variation of counter-move' yet chess is simple compared to the systems that scientists or politicians have to try to understand and predict in order to try to control. These themes of *uncertainty, nonlinearity, complexity and prediction* have been ubiquitous motifs of art, philosophy, and politics. We see them in Homer, where the gift of an apple causes the Trojan War; in Athenian tragedy, where a chance meeting at a crossroads settles the fate of Oedipus; in *Othello's* dropped handkerchief; and in *War and Peace* with Nikolai Rostov, playing cards with Dolohov, praying that one little card will turn out differently, save him from ruin, and allow him to go happily home to Natasha.[15]

Because of the 'unreasonable effectiveness of mathematics' in providing the 'language of nature' and foundations for a scientific civilization, we understand some systems very well and can make very precise predictions based on accurate quantitative models. Sometimes a mathematical model predicts phenomena which are later found (e.g. General Relativity's field equations); sometimes an experiment reveals a phenomenon that awaits an effective mathematical model (e.g. the delay between the discovery of superconductivity and a quantum theory). The work of mathematicians on 'pure' problems has often yielded ideas that have waited to be rediscovered by physicists. The work of Euclid, Apollonius and Archimedes on ellipses would be used centuries later by Kepler for his theory of planetary motion. The work of Riemann on non-Euclidean four-dimensional geometry was (thanks to Grossmann) used by Einstein for General Relativity. The work of various people since the 16th Century on complex numbers would be used by Heisenberg *et al* for quantum mechanics in the 1920s. The work of Cantor, Gödel, and Turing (c. 1860-1936) on the logical foundations of mathematics, perhaps the most abstract and esoteric subject, gave birth to computers. The work of Galois on 'groups' (motivated by problems with polynomial equations) would be used post-1945 to build the 'Standard Model' of particle physics using 'symmetry groups'. In a serendipitous 1972 meeting in the Institute of Advanced Study cafeteria, it was discovered that

---

[14] Divided by 2 because the journey ABC… is the same length as the reverse journey …CBA.

[15] 'I know that men are persuaded to go to war in one frame of mind and act when the time comes in another, and that their resolutions change with the changes of fortune… The movement of events is often as wayward and incomprehensible as the course of human thought; and this is why we ascribe to chance whatever belies our calculation.' Pericles to the Athenians.

the distribution of prime numbers has a still-mysterious connection with the energy levels of particles. G.H. Hardy famously wrote, in '*A Mathematician's Apology*' which influenced many future mathematicians, that the field of number theory was happily 'useless' and did not contribute to 'any warlike purpose'; even as he wrote the words, it was secretly being applied to cryptography and it now forms the basis of secure electronic communications among other things. Perhaps another example will be the 'Langlands Program' in pure mathematics which was developed in the 1960's and work on it is now funded by DARPA (the famous military technology developer) in the hope of practical applications. *The Millennium Problems are the most important outstanding problems in mathematics and experiments in large-scale collaborations to solve them should be a priority*, for their own sake and because of the (mostly unknowable) benefits that such breakthroughs bring.

Mathematicians invent (or discover?) concepts by abstraction and then discover connections between concepts.[16] Nature operates with universal laws[17] and displays symmetry[18] and regularity as well as irregularity and randomness.[19]

> 'What do we mean by "understanding" something? We can imagine that this complicated array of moving things which constitutes "the world" is something like a great chess game being played by the gods, and we are observers of the game. We do not know what the rules of the game are; all we are allowed to do is to watch the playing. Of course, if we watch long enough, we may eventually catch on to a few of the rules. *The rules of the game* are what we mean by *fundamental physics*. Even if we knew every rule, however, we might not be able to understand why a particular move is made in the game, merely because it is too complicated and our minds are limited. If you play chess you must know that it is easy to learn all the rules, and yet it is often very hard to select the best move or to understand why a player moves as he does. So it is in nature, only much more so; but we may be able at least to find all the rules. Actually, we do not have all the rules now. (Every once in a while something like castling is going on that we still do not understand.) Aside from not knowing all of the rules, what we really can explain in terms of those rules is very limited, because almost all situations are so enormously complicated that we cannot follow the plays of the game using the rules, much less tell what is going to happen next. We must, therefore, limit ourselves to the more basic question of the rules of the game. If we know the rules, we consider that we "understand" the world.' Richard Feynman.

These physical laws, or rules, use mathematicians' abstractions.[20]

---

[16] This happens in social sciences too. E.g. Brouwer's fixed-point theorem in topology was first applied to 'equilibrium' in economics by von Neumann (1930's), and this approach was copied by Arrow and Debreu in their 1954 paper that laid the foundation for modern 'general equilibrium theory' in economics (see below).

[17] '[M]any of the "macroscopic" laws of physics, such as the laws of thermodynamics or the equations of fluid motion, are quite universal in nature, with the microscopic structure of the material or fluid being studied being almost completely irrelevant, other than via some key parameters such as viscosity, compressibility, or entropy' (Tao).

[18] Symmetry principles are, says one of those who built the Standard Model of particle physics, 'like having a spy in the enemy's high command' (Weinberg), though like real spies they can also be misleading.

[19] 'Random' and 'pseudo-random' have precise technical meanings and are related to 'algorithmic (or Kolmogorov) complexity' (below).

[20] Einstein asked, 'How is it possible that mathematics, a product of human thought that is independent of experience, fits so excellently the objects of physical reality?' '*Is mathematics invented or discovered?*', Tim Gowers (Polkinghorne, 2011). Hilbert, Cantor and Einstein thought it is invented (formalism). Gödel thought it is discovered (Platonism). For a non-specialist summary of many issues concerning maths and prediction, cf. a talk by Fields Medallist Terry Tao. Wigner answered Einstein in a famous paper, '*The Unreasonable Effectiveness of Mathematics in the Natural Sciences*' (1960).

'It is an extraordinary feature of science that the most diverse, seemingly unrelated, phenomena can be described with the same mathematical tools. The same quadratic equation with which the ancients drew right angles to build their temples can be used today by a banker to calculate the yield to maturity of a new, two-year bond. The same techniques of calculus developed by Newton and Leibniz two centuries ago to study the orbits of Mars and Mercury can be used today by a civil engineer to calculate the maximum stress on a new bridge… But the variety of natural phenomena is boundless while, despite all appearances to the contrary, the number of really distinct mathematical concepts and tools at our disposal is surprisingly small… When we explore the vast realm of natural and human behavior, we find the most useful tools of measurement and calculation are based on surprisingly few basic ideas.' Mandelbrot

There is an extraordinary connection between mathematicians' aesthetic sense of 'beauty' and their success in finding solutions:

'It is efficient to look for beautiful solutions first and settle for ugly ones only as a last resort... It is a good rule of thumb that the more beautiful the guess, the more likely it is to survive.' Timothy Gowers.

'[S]ciences do not try to explain, they hardly even try to interpret, they mainly make models. By a model is meant a mathematical construct which, with the addition of certain verbal interpretations, describes observed phenomena. The justification of such a mathematical construct is solely and precisely that it is expected to work - that is, correctly to describe phenomena from a reasonably wide area. Furthermore, it must satisfy certain aesthetic criteria - that is, in relation to how much it describes, it must be rather simple… If only relatively little has been explained, one will absolutely insist that it should at least be done by very simple and direct means.' Von Neumann.[21]

Some of these models allow relatively *precise predictions* about a particular physical system: for example, Newton's equations for classical mechanics or the equations for 'quantum electrodynamics'.[22] Sometimes they are *statistical predictions* that do not say how a specific event will turn out but what can be expected over a large number of trials and with what degree of confidence: 'the epistemological value of probability theory is based on the fact that chance phenomena, considered collectively and on a grand scale, create a non-random regularity' (Kolmogorov). The use of statistical models has touched many fields: 'Moneyball' in baseball (the replacement of scouts' hunches by statistical prediction), predicting wine vintages and

---

[21] If possible, confirmation should come from areas which 'were not in the mind of anyone who invented the theory.' E.g. Newtonian mechanics was introduced to describe the planets' motion but it found application everywhere. Quantum Mechanics came into being because of problems in spectroscopy and connected areas, but it could predict things in chemistry and solid-state physics. One can have mathematically equivalent 'alternative interpretations of the same theory' and choose which one to use on the basis of 'mathematical convenience or taste'. E.g. Classical mechanics can be described in a 'causal' way, the Newtonian form, in which one describes the position and velocity of everything in a system then computes historical or future states. A second formulation is by 'the principle of minimum effect'. Similarly, one can choose between the Schrödinger or Heisenberg versions of quantum mechanics. Regardless of aesthetic or philosophical preferences, the one that will come to dominate is the one that 'lends itself better to formalistic extension towards valid new theories', and this is a 'formalistic, aesthetic criterion, with a highly opportunistic flavor.' Von Neumann.

[22] The equations of QED have been checked so delicately that, as Feynman said, 'If you were to measure the distance from Los Angeles to New York to this accuracy, it would be exact to the thickness of a human hair'.

ticket sales, dating, procurement decisions, legal judgements, parole decisions and so on (see Section 6 on Michael Nielsen).[23]

For example, many natural (e.g. height, IQ) and social (e.g. polling) phenomena follow the statistical theorem called the Central Limit Theorem (CLT) and produce a 'normal distribution', or 'bell curve'. Fields Medallist Terry Tao describes it:

> 'Roughly speaking, this theorem asserts that if one takes a statistic that is a combination of many independent and randomly fluctuating components, with no one component having a decisive influence on the whole, then that statistic will be approximately distributed according to a law called the normal distribution (or Gaussian distribution), and more popularly known as the *bell curve*...

> 'The law is universal because it holds regardless of exactly how the individual components fluctuate, or how many components there are (although the accuracy of the law improves when the number of components increases); it can be seen in a staggeringly diverse range of statistics, from the incidence rate of accidents, to the variation of height, weight, or other vital statistics amongst a species, to the financial gains or losses caused by chance, to the velocities of the component particles of a physical system. The size, width, location, and even the units of measurement of the distribution varies from statistic to statistic, but the bell curve shape can be discerned in all cases.

> 'This convergence arises not because of any "low-level" or "microscopic" connection between such diverse phenomena as car crashes, human height, trading profits, or stellar velocities, but because in all of these cases the "high-level" or "macroscopic" structure is the same, namely a compound statistic formed from a combination of the small influences of many independent factors. This is the essence of universality: the macroscopic behaviour of a large, complex system can be almost totally independent of its microscopic structure.

> 'The universal nature of the central limit theorem is tremendously useful in many industries, allowing them to manage what would otherwise be an intractably complex and chaotic system. With this theorem, insurers can manage the risk of, say, their car insurance policies, without having to know all the complicated details of how car crashes actually occur; astronomers can measure the size and location of distant galaxies, without having to solve the complicated equations of celestial mechanics; electrical engineers can predict the effect of noise and interference on electronic communications, without having to know exactly how this noise was generated; and so forth.'

Many other phenomena (e.g. terrorist attacks, earthquakes, stock market panics) produce a 'power law' and trusting to a *CLT model* of a phenomenon when it actually follows a power law causes trouble, as with the recent financial crisis.[24] When examining *phase transitions* of materials (e.g the

---

[23] In the 1950s people feared 'the hidden persuaders'. Now they fear data-mining. There are two main tools used regularly by organisations such as Amazon or Netflix: 1) 'regression analysis' (finding correlations in data) and 2) randomised trials (testing the effect of a particular thing (e.g. a drug, a political intervention) on a random population).

[24] When David Viniar (Goldman Sachs CFO) said of the financial crisis, 'We were seeing things that were 25-standard-deviation events, several days in a row,' he was implying that events that should happen once every 10^{135} years (the Universe is only $\sim 1.4 \times 10^{10}$ years old) were occurring 'several days in a row' - so he was either ignorant of basic statistics (unlikely) or not being honest (Buchanan, 2013). His statement refers to the probability of a '25-standard deviation event' happening *if* the underlying process follows the normal distribution. But stock prices are *fractal* and follow a *power law*, so large events like the 2008 crash are not actually as rare as predictions based on the normal distribution imply. Unfortunately many financial models are based on the normal distribution, hence they collapse in crises that they predict should not occur (see Section 7 and Endnote).

transition from water to ice), the patterns formed by atoms are almost always *fractals* which appear everywhere from charts of our heartbeats to stock prices to Bach.[25]

However, even our best understood mathematical models can quickly become practically overwhelming. Laplace voiced a famous expression of the post-Newton Enlightenment faith in science's potential to predict.

> 'We may regard the present state of the universe as the effect of its past and the cause of its future. An intellect which at a certain moment would know all the forces that animate nature, and all positions of the beings that compose it, if this intellect were vast enough to submit the data to analysis, would condense in a single formula the movements of the greatest bodies of the universe and those of the tiniest atom; for such an intellect nothing would be uncertain and the future just like the past would be present before its eyes... Present events are connected with preceding ones by a tie based upon the evident principle that a thing cannot occur without a cause that produces it... All events, even those which on account of their insignificance do not seem to follow the great laws of nature, are a result of it just as necessarily as the revolutions of the sun.' Laplace[26]

Newton himself had warned of the potential complexity of calculating more than two interacting bodies.

> 'The orbit of any one planet depends on the combined motions of all the planets, not to mention the action of all these on each other. But to consider simultaneously all these causes of motion and to define these motions by exact laws allowing of convenient calculation exceeds, unless I am mistaken, the force of the human intellect.'

It turned out that Newton's famous gravitational equation cannot be extended to just three bodies without producing '*deterministic chaos*',[27] so although 'cosmologists can use universal laws of fluid mechanics to describe the motion of entire galaxies, the motion of a single satellite under the influence of just three gravitational bodies can be far more complicated' (Tao). *Deterministic chaos*, a system which is 'sensitive to initial conditions', was first articulated by Poincaré as he struggled to solve the 'three-body problem' (chart below), and broke Laplace's dream of perfect understanding and prediction:

> 'If one seeks to visualize the pattern formed by these two [solution] curves and their infinite number of intersections, . . .[their] intersections form a kind of lattice-work, a weave, a chain-link network of infinitely fine mesh; ... One will be struck by the complexity of this figure, which I am not even attempting to draw. Nothing can give us a better idea of the intricacy of the three-body problem, and of all the problems of dynamics in general...

---

[25] 'Basic statistics such as the number of clusters, the average size of the clusters, or how often a cluster connects two given regions of space, appear to obey some specific universal laws, known as power laws... These laws seem to arise in almost every mathematical model that has been put forward to explain (continuous) phase transitions, and have also been observed many times in nature. As with other universal laws, the precise microscopic structure of the model or the material may affect some basic parameters, such as the phase transition temperature, but the underlying structure of the law is the same across all such models and materials.' Tao.

[26] A warning to academics: Napoleon was not impressed by Laplace and fired him after just six weeks saying, 'He was a worse than mediocre administrator who searched everywhere for subtleties, and brought into the affairs of government the spirit of the infinitely small.'

[27] The terms 'chaos' and 'complexity' have technical meanings different to normal use. A 'chaotic' system *appears* to be random but is actually *deterministic*; the apparent randomness comes from 'sensitive dependence' to feedback. Chaotic systems are a subset of nonlinear systems. 'Complexity' has many technical definitions such as 'algorithmic complexity' (how long is the shortest programme that outputs X?). Cf. Endnote.

'A very small cause which escapes our notice determines a considerable effect that we cannot fail to see, and then we say that that effect is due to chance. If we knew exactly the laws of nature and the situation of the universe at the initial moment, we could predict exactly the situation of that same universe at a succeeding moment. But even if it were the case that the natural laws had no longer any secret for us, we could still only know the initial situation *approximately*. If that enabled us to predict the succeeding situation *with the same approximation*, that is all we require, and we should say that the phenomenon had been predicted, that it is governed by laws. But it is not always so; it may happen that small differences in the initial conditions produce very great ones in the final phenomena. A small error in the former will produce an enormous error in the latter. Prediction becomes impossible, and we have the fortuitous phenomenon.' (Poincaré, *Science and Method*, 1913)

Box: Newton's deterministic equations for gravity (modern notation, Tao)

The equations of motion for the positions $x_1(t), …, x_N(t)$ of N particles of masses $m_1,…,m_N$ under Newtonian gravity is given by the system of equations

$$m_i \frac{d^2}{dt^2} x_i(t) = -\sum_{j \neq i} \frac{Gm_i m_j}{|x_j(t) - x_i(t)|^3} (x_j(t) - x_i(t))$$

for *i=1,…,N.*

Chart: For N=2 ('the two-body problem') these equations can be solved exactly (Tao)

orbit of B

Star A

m_A

focus

Barycenter

r_A

focus

r_B

orbit of A

Star B

m_B

$r = r_A + r_B$

For *N=2*, this system of equations can be solved exactly ("The two-body problem").

$M = m_A + m_B$

$m_A r_A = m_B r_B$

Mass of a Binary System
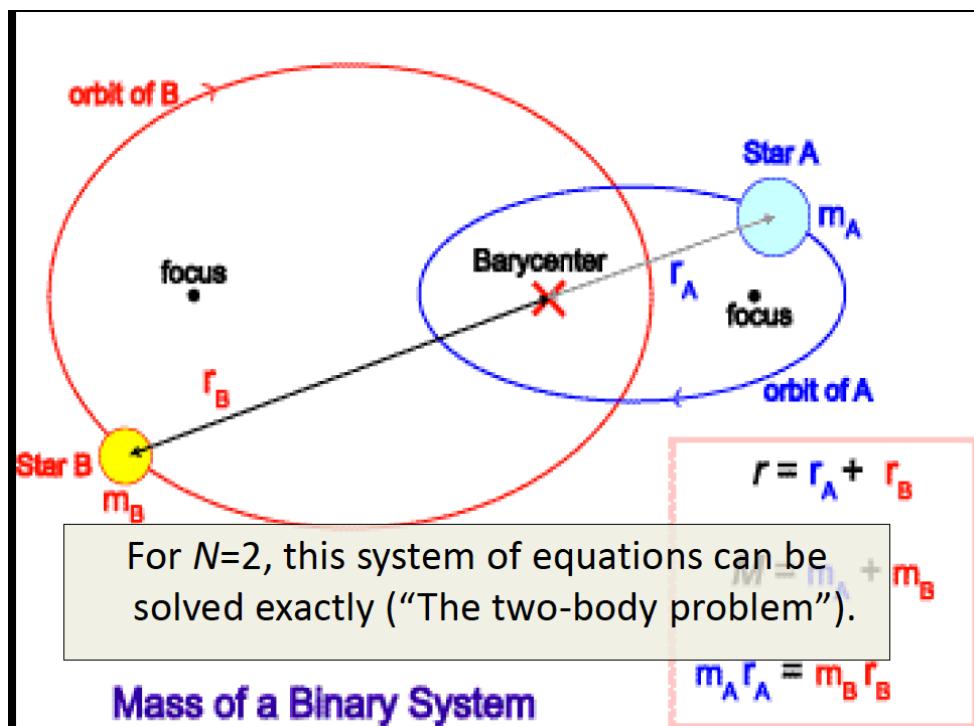
Chart: For N=3 ('the Three-Body Problem') the equations produce 'chaos' (Tao)

But even for *N*=3 - "the three-body problem" (the only problem to give Newton severe headaches) - there is no closed-form general solution to the equations.
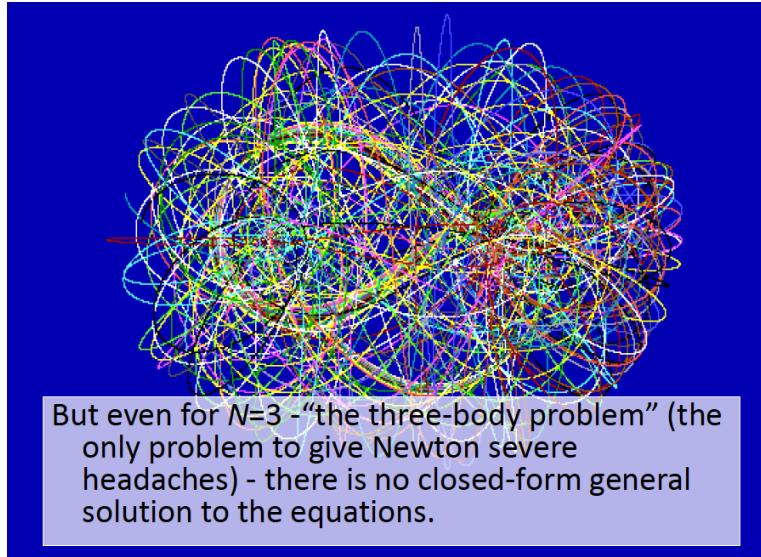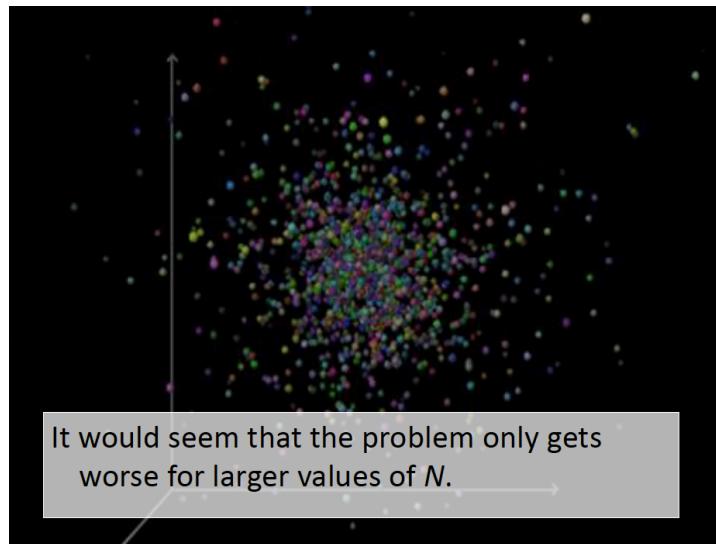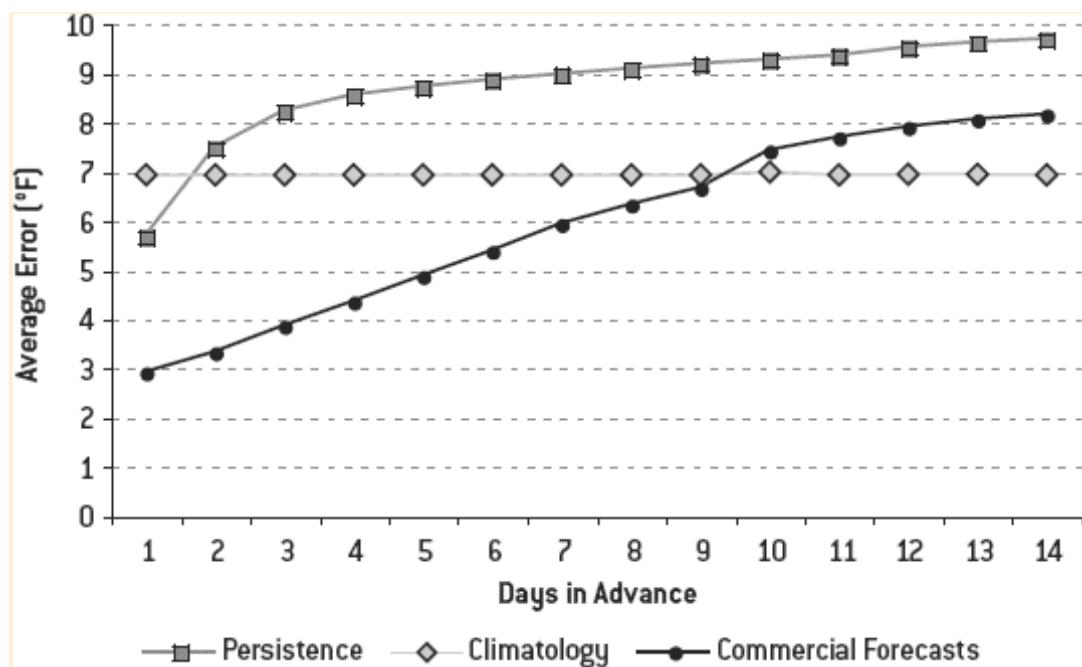
Chart: The problem gets worse as N increases (Tao)

It would seem that the problem only gets worse for larger values of *N*.

Even with systems displaying chaos because of sensitivity to initial conditions, short-term predictions are not hopeless. The best example is weather - the study of which was actually the prompt for Lorenz's re-discovery of 'chaos' (cf. Endnote). Weather forecasts have improved greatly over the past fifty years. For example, 25 years ago forecasts of where a hurricane would hit land in three days time missed by an average of 350 miles; now they miss by about 100 miles. We have bought ourselves an extra 48 hours to evacuate. Is a weather forecast better than it would be by simply a) looking at historical data (climatology), or b) assuming tomorrow will be similar to today (persistence)? Our forecasts are significantly better until about day 9 when forecasts become no better than looking at historical data.[28]

---

[28] Because of the great power of humans' visual cortex and image processing, they still add value to computer weather forecasts. The US Meteorological Prediction Center says that humans improve the accuracy of precipitation forecasts by about 25 percent over the computer guidance alone and temperature forecasts by about 10 percent. (Silver)

Chart: weather forecasts compared with climatology and persistence (Silver, 2012)

However, chaos means that beyond the short-term, forecasts rapidly break down and usually greater and greater resources are needed to extend the forecasts even just a little further; for example, there has been a huge increase in computer processing applied to weather forecasts since the 1950's, just to squeeze an accurate forecast out to Day 9.

> 'Even when universal laws do exist, it may still be practically impossible to use them to predict what happens next. For instance, we have universal laws for the motion of fluids, such as the Navier-Stokes equations, and these are certainly used all the time in such tasks as weather prediction, but these equations are so complex and unstable that even with the most powerful computers, we are still unable to accurately predict the weather more than a week or two into the future.' (Tao)

Between the precision of Newtonian mechanics (with a small number of interacting agents) and the statistics of multi-agent systems (such as thermodynamics and statistical mechanics) 'there is a substantial middle ground of systems that are too complex for fundamental analysis, but too simple to be universal. Plenty of room, in short, for all the complexities of life as we know it' (Tao).

Complex networks, emergent properties, and self-organisation

'If you took away the network structure of the human body, we would each be nothing more than a small box of chemicals and half a bathtub's worth of water.' Beinhocker.

Many of our most interesting problems can be considered as *networks*. Individual *nodes* (atoms, molecules, genes, cells, neurons, minds, organisms, organisations, computer agents, computer servers) and *links* (biochemical signals, synapses, internet routers, trade routes) form physical, mental, and cultural *networks* (molecules, cells, organisms, immune systems, minds, organisations, internet, biosphere, 'econosphere', cultures). Even a relatively simple network, with a small number of agents operating according to a small number of rules, can generate such vast complexity that an *exhaustive search* of all possible futures is unfeasible, such as chess (above). The most interesting

networks involve *feedback* (e.g. chemical signals, prices, neuronal firing, an assassination)[29] and are *nonlinear*, so one small event can have huge consequences: for example, turbulent physical systems or politics.[30] Complex networks have *emergent properties* including *self-organisation*: the relative strength of a knight in the centre of the chessboard is not specified in the rules but emerges from the nodes of the network (or 'agents') operating according to the rules.

The role of emergent properties in complex networks has many implications. In '*More Is Different*' (1972), one of the leaders of solid-state physics, Philip Anderson, argued that although scientists can agree on the 'reductionist hypothesis', the 'constructionist hypothesis' does *not* follow. The 'reductionist hypothesis' is that 'the workings of our minds and bodies, and all of the animate or inanimate matter of which we have any detailed knowledge, are assumed to be controlled by the same set of fundamental laws'. The 'constructionist hypothesis' is that the ability to reduce everything to simple fundamental laws implies 'the ability to start from those laws and reconstruct the universe' (and 'if everything obeys the same fundamental laws, then the only scientists who are studying anything really fundamental are those who are working on those laws').

> 'The constructionist hypothesis breaks down when confronted with the twin difficulties of scale and complexity. The behavior of large and complex aggregates of elementary particles … is not to be understood in terms of a simple extrapolation of the properties of a few particles. Instead, at each level of complexity entirely new properties appear.'

Although one can arrange sciences in a hierarchy from elementary particle physics through condensed matter physics, chemistry, molecular biology, cell biology etc, theoretically up to psychology and the social sciences (in the way E.O. Wilson describes in *Consilience*), 'this hierarchy does not imply that science X is "just applied Y." At each stage entirely new laws, concepts, and generalizations are necessary… Psychology is not applied biology, nor is biology applied chemistry.'

Ant colonies and *the immune system* are good examples of complex nonlinear systems with 'emergent properties' and self-organisation. The body cannot know in advance all the threats it will face so the immune system cannot be perfectly 'pre-designed'. How does it solve this problem? There is a large diverse population of individual white blood cells (millions produced per day) that sense threats (using different molecular 'receptors' created by random shuffling of DNA); if certain cells detect that a threat has passed a *threshold* (an immune response has costs and risks so the system uses thresholds), then they produce large numbers of daughter cells, with mutations, that are tested on captured antigens; unsuccessful daughter cells die while successful ones are despatched to the front; these daughter cells repeat the process so a rapid *evolutionary process* selects and reproduces the best defenders and continually improves performance; other specialist cells such as macrophages roam around looking for invaders that have been tagged by antibodies; some of the cells remain in the bloodstream, *storing information* about the attack, to guard against

---

[29] Eg. Alan Turing showed the importance of feedback in biology. *The Chemical Basis of Morphogenesis* (1952).

[30] 'We are poised on the tip of a lightning conductor, if we lose the balance I have been at pains to create, we shall find ourselves on the ground,' as Bismarck wrote to his wife during the 1871 peace negotiations. Online experiments have explored how complex social networks are subject to randomness that makes prediction extremely difficult at best because of the interdependence of decisions.

future attacks (immunity).[31] There is a constant evolutionary arms race against bacteria and other invaders.[32]

Looking at an individual ant, it would be hard to know that an ant colony is capable of farming, slavery, and war.

> 'The activity of an ant colony is totally defined by the activities and interactions of its constituent ants. Yet the colony exhibits a flexibility that goes far beyond the capabilities of its individual constituents. It is aware of and reacts to food, enemies, floods, and many other phenomena, over a large area; it reaches out over long distances to modify its surroundings in ways that benefit the colony; and it has a life-span orders of magnitude longer than that of its constituents… To understand the ant, we must understand how this persistent, adaptive organization emerges from the interactions of its numerous constituents.' (Hofstadter)[33]

Ant colonies face a similar problem to the immune system: they have to forage for food in an unknown environment with an effectively infinite number of possible ways to search for a solution. They send out agents looking for food; those that succeed return to the colony leaving a pheromone trail which is picked up by others and this trail strengthens. Decentralised decisions *via* interchange of chemical signals drive job-allocation (the division of labour) in the colony. Individual ants respond to the rate of what others are doing: if an ant finds a lot of foragers, it is more likely to start foraging.

Individual lymphocytes cannot access the whole picture; they *sample* their environment *via* their receptors. Individual ants cannot cannot access the whole picture; they *sample* their environment *via* their chemical processors. The receptor shape of lymphocytes and chemical processing abilities of ants are affected by *random* mutations; the way individual lymphocytes or ants respond has a random element. The individual elements (cells / ants) are programmed to respond *probabilistically* to new information based on the strength of signals they receive.

Environmental exploration by many individual agents coordinated *via* feedback signals allows a system to probe many different probabilities, reinforce success, 'learn' from failure (e.g withdraw resources from unproductive strategies), keep innovating (e.g novel cells are produced even amid a battle and ants continue to look for better options even after striking gold), and redundancy allows local failures without breaking the system at the macroscale. There is a balance between *exploring*

---

[31] How does the system avoid lymphocytes attacking healthy cells and the body itself? First, it seems that lymphocytes are tested in the bone marrow after birth before being released to see if they attack the body; if so, then they are not released. Second, regulatory T cells suppress the action of other T cells. Third, it is thought that perhaps cells that attack the body lose out in a competition for a chemical. When such measure fail, an autoimmune response is the result. How does it de-escalate the response if it is causing more damage than it is worth? There seem to be little-understood feedback signaling devices that limit responses. (Cf. Mitchell, 2009)

[32] Bacteria takeover cells' machinery and communications. They reprogram cells to take them over or trigger self-destruction. They disable immune cells and 'ride' them back into lymph nodes (Trojan horse style) where they attack. They shape-change fast so that immune cells cannot recognise them. They reprogram immune cells to commit suicide. They reduce competition by tricking immune cells into destroying other bacteria that help the body fight infection (e.g. by causing diarrhoea to flush out competition). (Finlay, *Scientific American,* 2010)

[33] Social and physical networks involve trade-offs between cost, efficiency, and fault tolerance. Making networks more robust means introducing redundant pathways that are not cost-effective in the short-term. Evolved biological networks provide clues about how designed networks could be improved. E.g. An experiment placed the slime mould *Physarum polycephalum* in space with 36 food sources representing the stations of the Tokyo area rail network, watched how it evolved a network, and compared it to the actual rail system. Such self-organised biological networks, with no central control or global information, can provide mathematical models that can improve designed systems such as remote sensor arrays, mobile *ad hoc* networks or wireless mesh networks. Cf. *Rules for Biologically Inspired Adaptive Network Design*, Tero et al (*Science*, 2010).

the immediate environment for information and *exploiting* that information to adapt. (Cf. Ch. 12 of Mitchell's *Complexity, 2009*)[34]

In such systems, *unintended consequences dominate*.[35] '*Tightly-coupled*' systems without 'safety gates', such as long domino chains without 'firebreaks', are not resilient in the face of nonlinear events (a gust of wind knocking over one domino can knock down the whole chain).[36] *Effects cascade*; 'they come not single spies but in battalions'. We are learning how *network topology* affects these dynamics. Many networks (including cells, brains, the internet, the economy) have a topology such that nodes are distributed according to a [power law](link) (not a bell curve), which means that the network looks like a set of *hubs and spokes* with a few spokes connecting hubs. This network topology (known technically as '[scale-free topology](link)') makes them *resilient* to random failure but *vulnerable* to the failure of crucial hubs that can cause destructive cascades (such as financial crises) - an example of the problems that come with nonlinearity.[37] Barabasi's group showed that similar topology and dynamics can be seen in networks operating at very different scales ranging from the tiny (e.g. 'Bose-Einstein condensates' at the quantum level) to the economy and internet. Many networks are *modular*, which allows specific modules to perform specific tasks, with links between them allowing broader coordination, and provides greater resilience to shocks.[38] Connections between *network topology*, *power laws* and *fractals* can be seen in work by physicist [Geoffrey West](link) both on [biology](link) and [cities](link) (see Endnote).

Adaptation to complexity

Complex systems are, therefore, hard to understand, predict, or control. Given our limited understanding of complex systems and their vast possibilities, success requires *adaptation*,[39] adaptation requires *prediction amid uncertainty*, and our evolved nature gives us not only amazing pattern recognition and problem-solving abilities but also *illusions*.[40] 'So little trouble do men take in the search after truth, so readily do they accept whatever comes first to hand' (Thucydides); 'men may construe things after their fashion / Clean from the purpose of the things themselves' (Cicero, *Julius Caesar*).[41] We often are governed by 'fear, honour and interest' (Thucydides), we attribute success to skill and failure to luck,[42] and we prefer to enhance prestige rather than admit ignorance or

---

[34] There are costs of decentralisation too. E.g. The decentralised Kmart was destroyed by WalMart's use of computers to create a network of shops operating in a more centralised way.

[35] '[T]hat which in the first instance is prejudicial may be excellent in its remoter operation; and its excellence may arise even from the ill effects it produces in the beginning. The reverse also happens: and very plausible schemes, with very pleasing commencements, have often shameful and lamentable conclusions' (Burke).

[36] Cf. *Normal Accidents*, Charles Perrow.

[37] E.g a disgraced or assassinated leader, a vital unit watching porn on the web, like the infamous SEC regulators. Cf. Barabasi's *Linked*.

[38] E.g. Chongqing (China) saw the evolution of a new ecosystem for designing and building motorbikes in which 'assembler' companies assemble modular parts built by competing companies, instead of relying on high quality vertically integrated companies like Yamaha. This rapidly decimated Japanese competition.

[39] In one of the most famous speeches, Pericles singled out the Athenian quality of *adaptation* (literally 'well-turning') as central to its extraordinary cultural, political and economic success.

[40] Our brains evolved to solve social and practical problems, not to solve formal mathematical problems. Translating mathematical and logical problems into social problems makes them easier for people to solve. Cf. [Nielsen](link).

[41] 'People are not accustomed to thinking hard and are often content to trust a plausible judgment that comes quickly to mind.' Kahneman. But this is not *necessarily* foolish (cf. Vernon Smith, section 7).

[42] 'The advantage scientists bring into the game [trading] is less their mathematical or computational skills than their ability to think scientifically. They are less likely to accept an apparent winning strategy that might be a mere statistical fluke' (Jim Simons, mathematician and founder and CEO of Renaissance, the most successful hedge fund).

error and face reality. Accurate prediction is therefore hard[43] and we struggle to adapt in a 'fog of war' that makes the simplest things hard and failure the norm.[44] Faced with such complexity, politicians and others have operated mostly on heuristics, guesswork, willpower and tactical adaptation (see Section 6 below).

Quantitative models, design, and engineering

'Because we understand the universal physical laws that govern matter and energy, we understand the physical laws that will govern the material structures of future technologies… [G]rowing computational capacity can enable us to simulate systems that have not yet been built. New aircraft typically fly as expected, new computer chips typically operate as expected. These same capabilities can also be used to simulate systems that cannot yet be built… Using computational simulation this way is like the earlier use of telescopes to view planets that spacecraft could not yet reach.' Drexler

Aeronautics was confined to *qualitative stories* (like Icarus) until the 1880s when people started making careful *observations and experiments* about the principles of flight. *Empirical design rules* were refined. Decades later we developed *quantitative models* and these moved to computers. Each part had its own computational model, or '*device physics*', and the overall design is a synthesis of the device physics of all the components. The Boeing 747 is built from ~50,000 kinds of parts, each with their own device physics, and ~6m total components. Now theory is so accurate that airplanes are designed and tested on computers and built mainly without wind-tunnel testing. The models are so accurate that autopilots based on them are relied on and they can be used to design and fly planes, such as the X-29, that are too unstable for humans to fly unassisted by computers. Similarly nuclear weapons are so well understood that they can be designed and built from simulations.

There are some assumptions: 1) all the parts come from careful design, 2) the parts function according to the design, 3) the parts can be reliably put together and *the whole behaves as expected* (i.e it is *not* like an ant colony or immune system). Mature engineering disciplines, such as aeronautics or electronics, now rely on a) *standard test and measurement gear* (oscilloscopes, pressure gauges etc), and b) *computer-aided design* (CAD) software, like SolidWorks (mechanical engineering) or Spice and Verilog (circuits), that are based on *quantitative models* of the device physics of all the components and '*computational fluid dynamics*' software, which yield precise predictions and which can send instructions to automated manufacturing tools. There *are models, algorithms, and computers that bring precise understanding, prediction, and control* (cf. *Biology is Technology*, Carlson).

We are extending this approach to biology. X-ray crystallography and nuclear magnetic resonance gave us access to the molecular world and our abilities to see molecules and engineer computers are merging in biology and creating new fields such as [synthetic biology](#) and *computational biology* (see Section 4). We are beginning to apply the principles of mature engineering disciplines but we have very few quantitative models of biology or design tools. Our summaries of biological networks generally involve the names of molecules and arrows showing interactions, with very few numbers. The *Mycoplasma genitalium*, a tiny bacterium, operates with only ~430 genes but even so we still have not found a function for ~100 of its essential proteins. Yeast consists of ~6,300 kinds of genetic parts: we have low resolution models for ~50% of them and the quantitative device physics for only a handful. The total number of moving parts in a yeast cell is unknown but is 'many, many

---

[43] Cf. Tetlock's '*Expert Political Judgement*' on how predictions by 'political experts' are barely better than chimps throwing darts at a board, a book that ought to be much-studied in Westminster and Whitehall but is practically unknown. Klein's '*Sources of Power*' is interesting on how experts really make decisions under pressure. Cf. Section 6.

[44] Cf. Endnote: Constant Failure.

millions'. The human genome encodes ~21,000 different proteins. We understand the function of only a fraction of these. We are far from having precise quantitative models of all the molecular interactions in a cell though we are making rapid progress and there is hope that we can 'bootstrap' our way to superior understanding by mapping progressively more complex organisms.. (Carlson, *ibid*, and see Section 4)

This approach is also being extended to dealing with traditional problems of manufacturing failure. Because statistically accurate test-to-failure simulations of physical objects are too expensive for anything other than the cheapest things (you need thousands of tests to get good data on a 1-in-thousands failure), other methods are needed. Ideally, one uses computer simulations as they are so cheap. E.g. you import a CAD file of the part into statistical stress-modelling software like Finite Element Analysis (a programme that simulates pressure, temperature etc effects on the CAD model). However, accurately predicting actual material behaviour under stress has been impossible because we do not understand exactly how materials crack at the micromaterial level - tiny variations in the raw material and manufacturing lead to very different performance, so real products fail in a normal distribution and not all at the same time. Vextec therefore has created *Virtual Life Management*, a tool that takes real detail on the microstructure of a material and how it varies in real life, creates a large number of virtual samples varying according to the real world data, then applies stress to them using the information from Finite Element Analysis. It has already had some dramatic successes; if it lives up to expectations, then it and its successors will revolutionise the field and save lives and billions of dollars. (Cf. *Why things fail*, Wired 2013.)

This combination of a) precise real world data, b) mathematical modelling software that captures dynamic processes, and c) the potential for simulations of large numbers of virtual samples that vary according to real world data, has the potential to revolutionise our understanding, prediction, and control of complex systems and improve how we cope with failure.

However, there is also a danger in the use of statistical models based on 'big data' analysis - 'overfitting' models and wrongly inferring a 'signal' from what is actually 'noise'. We usually a) have a noisy data set and b) an inadequate theoretical understanding of the system, so we do not know how accurately the data represents some underlying structure (if there is such a structure). We have to infer a structure despite these two problems. It is easy in these circumstances to 'overfit' a model - to make it twist and turn to fit more of the data than we should, but then we are fitting it not to the signal but to the noise. 'Overfit' models can seem to explain more of the variance in the data - *but they do this by fitting noise rather than signal* (Silver, *op. cit*).

This error is seen repeatedly in forecasting, and can afflict even famous scientists. For example, Freeman Dyson tells a short tale about how, in 1953, he trekked to Chicago to show Fermi the results of a new physics model for the strong nuclear force. Fermi dismissed his idea immediately as having neither 'a clear physical picture of the process that you are calculating' nor 'a precise and self-consistent mathematical formalism'. When Dyson pointed to the success of his model, Fermi quoted von Neumann, 'With four parameters I can fit an elephant, and with five I can make him wiggle his trunk', thus saving Dyson from wasting years on a wrong theory (*A meeting with Enrico Fermi*, by Freeman Dyson). Imagine how often people who think they have a useful model in areas not nearly as well-understood as nuclear physics lack a Fermi to examine it carefully.

Can we extend such methods to the understanding, prediction and control of human behaviour?

It has often been argued (e.g by Kant[45]) that the 'complexity' of human behaviour renders precise mathematical treatment impossible, or that the undoubted errors of modern economics in applying the tools of mathematical physics, are evidence of the irredeemable hopelessness of the goal. However, in the opening of his '*Theory of Games and Economic Behaviour*' (which created the field of 'game theory' and was one of the most influential books ever written on economics), written hastily during breaks from pressing wartime business such as the Manhattan Project, the famous mathematician John von Neumann explained the real problems in applying mathematics to society or culture and why Kant was wrong. [46]

'It is not that there exists any fundamental reason why mathematics should not be used in economics. The arguments often heard that because of the human element, of the psychological factors etc., or because there is – allegedly – no measurement of important factors, mathematics will find no application, can all be dismissed as utterly mistaken. Almost all these objections have been made, or might have been made, many centuries ago in fields where mathematics is now the chief instrument of analysis [e.g. physics in the 16th Century or chemistry and biology in the 18th]…

'As to the lack of measurement of the most important factors, the example of the theory of heat is most instructive; before the development of the mathematical theory the possibilities of quantitative measurements were less favorable there than they are now in economics. The precise measurements of the quantity and quality of heat (energy and temperature) were the outcome and not the antecedents of the mathematical theory...

'The reason why mathematics has not been more successful in economics must be found elsewhere… To begin with, the economic problems were not formulated clearly and are often stated in such vague terms as to make mathematical treatment *a priori* appear hopeless because it is quite uncertain what the problems really are. *There is no point using exact methods where there is no clarity in the concepts and issues to which they are applied.* [Emphasis added] Consequently the initial task is to clarify the knowledge of the matter by further careful descriptive work. But even in those parts of economics where the descriptive problem has been handled more satisfactorily, mathematical tools have seldom been used appropriately. They were either inadequately handled … or they led to mere translations from a literary form of expression into symbols…

'Next, the empirical background of economic science is definitely inadequate. Our knowledge of the relevant facts of economics is incomparably smaller than that commanded in physics at the time when mathematization of that subject was achieved. Indeed, the decisive break which came in physics in the seventeenth century … was possible only because of previous developments in astronomy. It was backed by several millennia of systematic, scientific, astronomical observation, culminating in an observer of unparalleled calibre, Tycho de Brahe.

---

[45] 'For it is quite certain that in terms of merely mechanical principles of nature we cannot even adequately become familiar with, much less explain, organized beings and how they are internally possible. So certain is this that we may boldly state that it is absurd for human beings even to attempt it, or to hope that perhaps some day another Newton might arise who would explain to us, in terms of natural laws unordered by any intention, how even a mere blade of grass is produced' (*Critique of Judgement*).

[46] Von Neumann (documentary) was one of the leading mathematicians of the 20th Century, was a major contributor to the mathematisation of quantum mechanics, created the fields of 'quantum logic' (1936) and 'game theory' (during World War II), and (with Turing) was one of the two most important creators of modern computer science (including AI) which he developed partly for immediate problems he was working on (e.g. the hydrogen bomb and ICBMs) and partly to probe the general field of understanding complex nonlinear systems. Cf. Endnote.

Nothing of this sort has occurred in economics. It would have been absurd in physics to expect Kepler and Newton without Tycho – and there is no reason to hope for an easier development in economics…

'Very frequently the proofs [in economics] are lacking because a mathematical treatment has been attempted in fields which are so vast and so complicated that for a long time to come – until much more empirical knowledge is acquired – there is hardly any reason at all to expect progress *more mathematico*. The fact that these fields have been attacked in this way … indicates how much the attendant difficulties are being underestimated. They are enormous and we are now in no way equipped for them.

'[We will need] changes in mathematical technique – in fact, in mathematics itself… It must not be forgotten that these changes may be very considerable. The decisive phase of the application of mathematics to physics – Newton's creation of a rational discipline of mechanics – brought about, and can hardly be separated from, the discovery of the infinitesimal calculus… The importance of the social phenomena, the wealth and multiplicity of their manifestations, and the complexity of their structure, are at least equal to those in physics. It is therefore to be expected – or feared – that mathematical discoveries of a stature comparable to that of calculus will be needed in order to produce decisive success in this field… *A fortiori*, it is unlikely that a mere repetition of the tricks which served us so well in physics will do for the social phenomena too. The probability is very slim indeed, since … we encounter in our discussions some mathematical problems which are quite different from those which occur in physical science.'

Von Neumann therefore exhorted economists to humility and the task of 'careful, patient description', a 'task of vast proportions'. He stressed that economics could not attack the 'big' questions - much more modesty is needed to establish an exact theory for very simple problems, and build on those foundations.

'The everyday work of the research physicist is … concerned with special problems which are "mature"… Unifications of fields which were formerly divided and far apart may alternate with this type of work. However, such fortunate occurrences are rare and happen only after each field has been thoroughly explored. Considering the fact that economics is much more difficult, much less understood, and undoubtedly in a much earlier stage of its evolution as a science than physics, one should clearly not expect more than a development of the above type in economics either...

'The great progress in every science came when, in the study of problems which were modest as compared with ultimate aims, methods were developed which could be extended further and further. The free fall is a very trivial physical example, but it was the study of this exceedingly simple fact and its comparison with astronomical material which brought forth mechanics. It seems to us that the same standard of modesty should be applied in economics… The sound procedure is to obtain first utmost precision and mastery in a limited field, and then to proceed to another, somewhat wider one, and so on.'

Von Neumann therefore begins with 'the behavior of the individual and the simplest forms of exchange' with the hope that this can be extended to more complex situations.

'Economists frequently point to much larger, more 'burning' questions… The experience of … physics indicates that this impatience merely delays progress, including that of the

treatment of the 'burning' questions. There is no reason to assume the existence of shortcuts...

'It is a well-known phenomenon in many branches of the exact and physical sciences that very great numbers are often easier to handle than those of medium size. An almost exact theory of a gas, containing about 10^{25} freely moving particles, is incomparably easier than that of the solar system, made up of 9 major bodies… This is … due to the excellent possibility of applying the laws of statistics and probabilities in the first case.

'This analogy, however, is far from perfect for our problem. The theory of mechanics for 2,3,4, … bodies is well known, and in its general theoretical …. form is the foundation of the statistical theory for great numbers. For the social exchange economy – i.e. for the equivalent 'games of strategy' – the theory of 2,3,4… participants was heretofore lacking. It is this need that … our subsequent investigations will endeavor to satisfy. In other words, only after the theory for moderate numbers of participants has been satisfactorily developed will it be possible to decide whether extremely great numbers of participants simplify the situation.' (See below for more on von Neumann and economics.)

While some of von Neumann's ideas were extremely influential on economics, his general warning here about the right approach to the use of mathematics was not heeded. However, outside orthodox economics and financial theory (the failures of which were made obvious by the 2008 financial crisis) we are making progress with new ideas and tools such as automated reasoning, agent-based models (ABMs), and network dynamic analysis, described below (Section 7 and Endnote).[47]

Can we make breakthroughs in *understanding, prediction and control* of complex nonlinear systems that will help us deal with future shocks such as new conflicts, enterprises, technologies, disasters, and discoveries - breakthroughs that would make current systems seem as archaic as Ptolemy's epicycles did after Newton? How many breakthroughs will be secret? For example, given that the NRO can give NASA Hubble-like cast-offs, what is it now using secretly?[48] Cf. *'Collaborative Mathematical Innovation Challenges'* by *Think Tank Maths* for a detailed research agenda.

*

---

[47] Cf. this paper by Andrew Lo on differences between economics and physics and a proposed 'taxonomy' of uncertainty comparing different types of physics experiment with economics, ranging from deterministic Newtonian mechanics, through noisy systems and statistical predictions to 'irreducible uncertainty'. Economics remains in many ways in the state von Neumann lamented; even the basic facts are controversial. Cf. Endnote.

[48] E.g. Renaissance Technologies, a very successful and secretive hedge fund, hires leading pattern-recognition experts. The Heilbronn Institute's partnership with GCHQ on applied mathematics: how many other breakthroughs do GCHQ, NSA *et al* monitor? IARPA (the intelligence community's version of DARPA) has its FUSE programme aimed at identifying emerging technologies and OSI programme to predict social events, and is likely to be thinking about how to spy on the likes of Renaissance.

2. Energy and space. *How soon can we develop solar energy to provide limitless clean energy? Can we open up commercial development of the solar system to monitor and safeguard the earth, develop new industries, and propel further scientific discoveries?*

Energy[49]

Some basic facts… Humans consume about 14 terrawatts (1.4×10^{13} W) of *energy* in total; global *power* consumption is about 2.5 terrawatts. US total *energy* consumption is ~3.5 TW, or ~10^4 watts per person, which can be pictured as a flow of ~300 tons per second of oil; US *power* consumption is ~400 GW. UK total *energy* consumption is ~300 GW and *power* consumption is ~40GW. Compared to the ~14TW consumed by human civilization, the sun provides the earth with ~170,000 TW, which is roughly 10,000 times more than we use, so the sun provides in an hour what we need for a year.

One standard nuclear or coal power station (or a square kilometer of sunlight hitting earth) produces about 1 GW, so ~1/3,500[th] of US needs. Average power consumption for a US house is ~1-2 KW. New York uses ~10 GW but China has been installing 1GW of coal power stations *every week*. China's oil use has been growing by ~50% per year. Each 10% cut in US carbon emissions is 'completely negated by six months of China's emission growth.' US cars are expected to produce a rise in global temperature of only *one-fortieth of a degree* in the next 50 years; our real problems lie elsewhere (Muller).

Oil demand is projected to grow from ~80m barrels per day to ~120m per day by 2030 but the rate of discovery of new conventional supplies has plateaued and the biggest known fields are producing less. Many predictions of 'peak oil' have been wrong. In 2010 the International Energy Agency began warning of the need to prepare for declining oil production while simultaneously predicting that *worldwide energy demand will double by 2050*. However, vast new discoveries of 'shale oil' and 'shale gas' have recently been made.

Increasing energy productivity… We have significantly improved energy productivity. Already in the US and Europe energy use per capita has stabilised (the number of barrels of oil per capita has fallen significantly). As the rest of the world becomes richer and invests in technology (benefiting from the developed world's R&D), it will become more energy efficient and use technology better to minimise environmental impact, as we have.[50]

Even small improvements in energy efficiency sustained over decades will bring great advances.
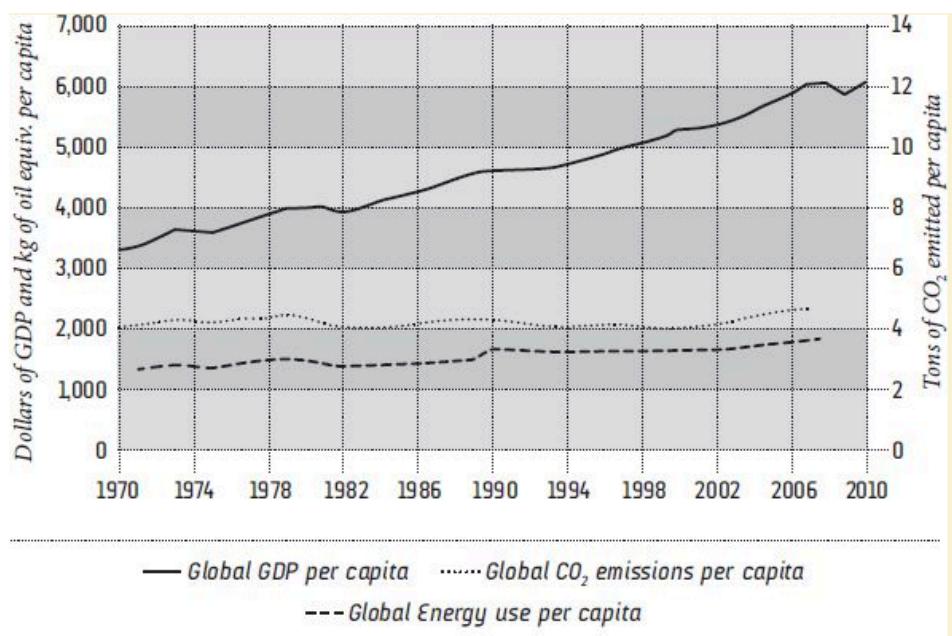
> 'From 1845 to the present, the amount of energy required to produce the same amount of gross national product has steadily decreased at the rate of about 1 percent per year... One percent per year yields a factor of 2.7 when compounded over 100 years. It took 56 BTUs (59,000 joules) of energy consumption to produce one (1992) dollar of GNP in 1845. By 1998, the same dollar required only 12.5 BTUs (13,200 joules)… [A]t 2 percent growth ... conservation outruns population by a large factor. Two percent compounded over 100 years reduces energy use by a factor of 7.2. By 2100, with a world population of 10 billion people, everyone can be living at the current European standard of living and yet expending half the energy we are using today.' (Muller)

---

[49] Much of this section comes from *Energy for Future Presidents* (2012) by Richard Muller, Professor of Physics at Berkeley and an expert on energy policy.

[50] E.g '… lighting has progressed from candles producing one lumen for every six watts of candle wax burned, to incandescent electric bulbs producing 17 lumens with each watt of electric input, to fluorescent lighting that produces 100 lumens per watt' (Rosenfeld).

Graph: increasing energy productivity (Naam)

Cost of electricity… The cost of energy varies widely according to source. The typical cost of electricity bought from a US utility is ~10 cents per kilowatt-hour (kWh). Petrol costs about 2.5 times more than retail natural gas per unit of energy, and seven times more than wholesale gas. The energy from a AAA battery costs about 10,000 times more than electricity from the wall. Depending on different assumptions and details, a new gas power station produces energy at ~6-7 cents per kilowatt-hour (kWh); coal ~10 cents; nuclear ~11-12 cents; wind varies enormously between ~10-25 cents (much more expensive offshore); solar ~20 cents. (Muller, p. 142)

Shale gas… The main competitor to nuclear, wind, and solar power is natural gas. The discovery of vast reserves of shale gas and the technology to extract it relatively cheaply is changing thinking about energy across the world (and on the ocean floors there is many times more deposits of frozen methane). Proven US natural gas reserves have risen from ~200 trillion cubic feet (2001) to ~300 trillion (2010) and ~860 trillion (2011). This number may rise to thousands of trillions of cubic feet. For equal energy, gas produces about half the CO^2 of coal.
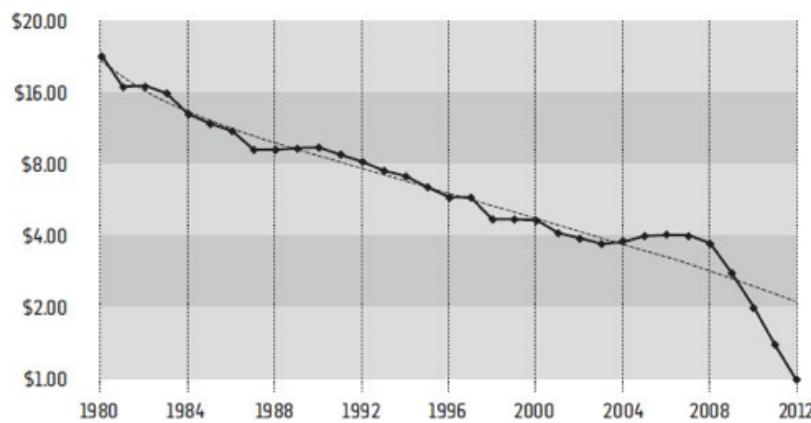
Shale oil… Shale oil is like shale gas - trapped in rock and hitherto uneconomic to exploit. The technology has suddenly changed the landscape in the last two years: the energy of the extracted oil is 3.5 times larger than the energy used to extract it. This potentially provides another windfall in the form of trillions of barrels of shale oil. Muller estimates that 25% of US oil may come from shale oil by 2020 and that much of it will be extracted for as little as $30 per barrel. This will enormously help countries become independent of the Middle East but might also cause large environmental harm.

Synfuel… Synfuel (coal or gas turned to oil) can be produced for ~$60 per barrel now. 'Synfuel manufactured from natural gas is one of the few energy sources that might beat out compressed natural gas.' Shale oil could be cheaper to produce than synfuel.

Solar… Sunlight delivers ~1KW (10^3 watts) of power per square metre[51] and ~1GW (10^9 watts) per square kilometer. The best solar cells convert ~40% of that energy to electricity, so a square mile of sunlight converted at ~40% yields over 1GW - about the same as a large coal or nuclear plant.

Using today's solar technology, the amount of land needed to provide double our current global energy consumption is only about one-sixteenth of the world's deserts. 'The Sahara Desert, alone, is more than five times the area we would need to capture enough energy to power humanity in 2050, using 2010 consumer-grade technology.' In 1980, average US retail electricity cost ~5 cents per KWh (today's dollars) but electricity produced from solar power cost nearly $20. While retail electricity is now ~10 cents per KWh, solar costs have plummeted towards parity; they are not there yet but will be soon. The first solar photovoltaic panel built by Bell Labs in 1954 cost $1,000 per watt of electricity; the cost in 1980 was ~$20 per watt; in 2012 the cost fell below $1 per watt (2012 dollars). The fall in cost is accelerating. (Naam, 2013).

Graph: the cost of electricity from solar photovoltaic cells is ~1/20th of the 1980 price (Naam)

As the number of installed cells grows, the price continues to fall as expected efficiency savings from growing understanding kick in (a 'learning curve'): for every doubling of the amount of solar energy built, prices drop around 20 percent (a much higher learning curve than oil and gas's 3-4%). Solar photovoltaic cells manufactured in 2010 capture as much energy as it took to build and install them within the first three years of their estimated thirty-year lifetime giving an 'energy return on investment' (EROI) of 10 and this is improving fast. If the energy payback time falls to just one year, the EROI will be 30 - about the same level as natural gas and oil. (Naam, 2013)

However, the standard figures often used for solar involve *peak* watts but there is cloud, night, and other things that degrade performance. A solar cell delivers on average only about ~1/8 of its peak performance. The largest manufacturer of solar cells in the world is China's Suntech which installed about 1 GW in 2011 but this equates to a real output of only ~1/8 GW. *This compares with China's 50 GW of installed coal power - about 400 times more than Suntech's solar cells.* Further, even if the cost of solar cells falls to a negligible amount, as many think it will, there is still a cost of installation and maintenance (and problems of transporting the electricity). Muller concludes that solar may not be competitive with gas in the West even if its cost continues to fall quickly but that it may well be competitive in the developing world where installation and maintenance are cheaper.[52]

---

[51] 2m^2 of sunlight on the roof of a car converted at 40% gives ~800 watts so solar powered cars are not viable.
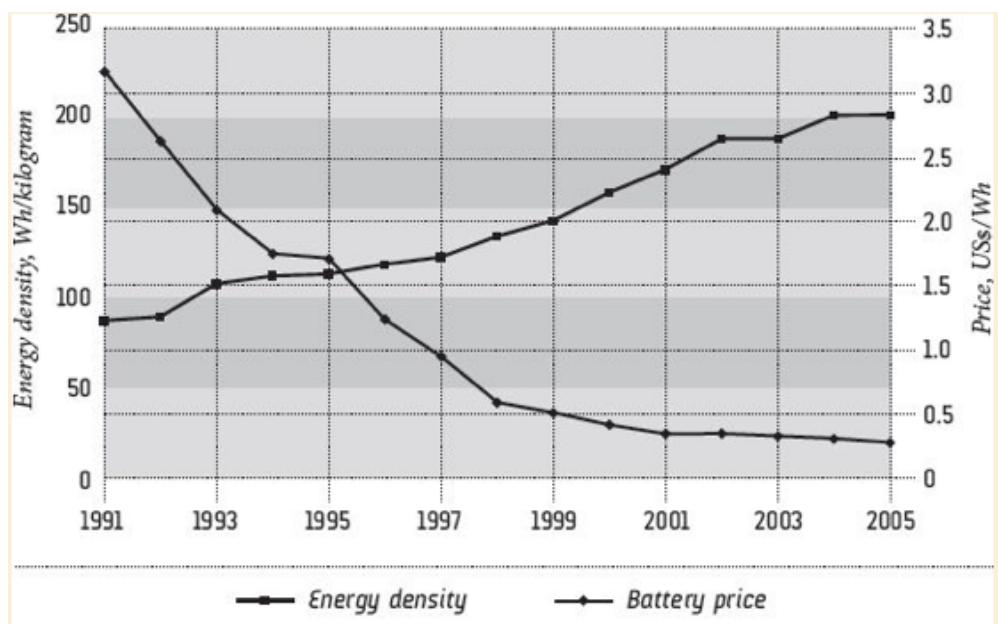
[52] Solar-thermal power plants are unlikely to be cost-effective.

Nuclear… 'Nuclear power is safe, and waste storage is not a difficult problem. Fears are driven by unfamiliarity and misinformation.' Disasters such as Fukushima were 'not nearly as catastrophic as many people think, and they should not imply any major change in energy policy' (Muller).[53] Small modern modular nuclear power plants can be made much more efficient and safe. (Older nuclear reactors like Fukushima work on 'active safety': i.e people have to keep doing things to keep it safe and if they stop then things go wrong. New nuclear reactors almost all work on 'passive safety' ('walk away safe'): i.e if people forget, leave, die or otherwise stop acting, the plant will continue to be safe.) China is building or planning to build ~75 new nuclear plants with another 100 proposed. Even if the West and Japan ignore nuclear, the rest of the world will not. In America, nearly half of rail shipping by weight is coal being moved to power plants; you avoid shipping over 500,000 tons of coal by shipping just one ton of enriched uranium.

Although there have been hopes since the successful testing of nuclear *fusion* bombs that we would solve our energy needs with fusion (which uses hydrogen, the most abundant material, for fuel), it remains elusive. Muller thinks that the problems will be solved this century but it looks at least twenty years away and has been so for over half a century (cf. *Scientific American,* March 2010, for a survey of some of the problems).

Batteries… Between 1991-2005, the price of storing energy in a battery fell by a factor of ~10 while the amount of energy that can be stored per kg has more than doubled. This improvement is better than for solar so if it continues then our ability to store solar anergy will improve faster than our ability to collect it.

Graph: rising energy density and falling cost of stored energy in batteries (1990-2005)

Electric and hybrid cars… However, this improvement is unlikely to make all-electric cars viable unless it is much more dramatic than seems likely. The performance of electric cars such as the Tesla Roadster is impressive. The problem is its batteries weigh ~1,000 pounds or nine times as much as the car's engine, cost ~$40,000 per car, and need to be replaced after a few hundred charges. In comparison, the petrol in a fifteen-gallon automobile tank weighs around 100 pounds. Lead-acid car batteries are very efficient (we get out 80-90% of the energy we put in). However,

---

[53] The radioactivity near Fukushima 'dropped within months to a level lower than that normally experienced in Denver' (Muller).

the energy stored per gram is small compared to that in normal fuels: petrol provides ~250 times more energy per gram than a car battery so even after factoring in a conversion rate of only 20% (in a generator turning the petrol to electricity) petrol still provides 50 times the electric energy of an equal-weight battery. Further, electric cars that derive their electricity from coal are *worse* for the environment than petrol cars. Muller expects the interest in all-electric cars 'to die soon, when the batteries need replacement and the owners discover how much they cost.'[54]

However, he expects most of us to be driving hybrids made of lighter, stronger materials 'within the next decade or two'.

> 'Hybrid autos have a great future, but plug-in hybrids and all-electric automobiles do not: they cost much more to operate than do gasoline cars once you include battery replacement costs … [with the exception of cars] that run on lead-acid batteries, with very short ranges (on the order of 40-60 miles) [which] could achieve widespread use in China, India, and the rest of the developing world'].' (Muller)

Discussion of electric and hybrid cars should not ignore the tiny effect of US car emissions on CO_2 levels - just 1/40th of a degree over decades. This is much less of a problem than most media coverage implies.

Hydrogen, tidal, geothermal, waves, biofuels...

According to Muller, 'The hydrogen economy is going nowhere.' Extracting hydrogen *requires* energy and when we use hydrogen as fuel we get only a fraction of that energy back out. It takes up a large volume so contains much less energy *per litre* than petrol (it contains more energy per gram). Hydrogen as a fuel costs ~25 cents per kWh - much more expensive than natural gas which 'can store equal energy in a quarter of the space that hydrogen requires.'

'[G]eothermal, tidal energy, and wave energy also have no large scale future.' Corn ethanol should not qualify as a 'biofuel' and its use does not reduce carbon emissions. Ethanol from cellulose has some hope but Muller thinks it will have a tiny effect on global warming and 'come along too late, and be too expensive, to truly compete with compressed natural gas, synfuel, or shale gas.'

A UK ARPA-E… The UK ought to establish an equivalent to America's ARPA-E (modelled on DARPA), to pursue high-risk/high-payoff energy projects in strategic areas - and make sure that, similarly to DARPA, it operates outside all normal bureaucratic systems (e.g. Whitehall personnel systems, EU procurement rules) and has the same philosophy: focus on things that are *not* incremental improvements, do not operate your own labs but piggy-back on Universities, give people space to innovate, and close them down if they do not deliver in a 3-5 year period.

Space

Viking trips to America did not change the world; opening the Atlantic to commerce changed the world. Various attempts to reduce drastically the cost of getting materials into orbit have failed, including NASA's and DARPA's recent attempts. [SKYLON](), a hypersonic spaceplane developed from a classified UK 1980's programme, might open up space by dealing with the problem that heat increases as speed increases, rendering hypersonic spaceplanes *hitherto* impractical. 'Skylon's technology would cut the cost of launching a satellite from around £15,000 per kg to around £650 per kg' (i.e.. a thirtieth). The UK Government and European Space Agency regard Skylon's engine as

---

[54] Two possibilities with promise for improving car productivity are a) lighter, stronger materials and b) 'regenerative braking', in which the energy from braking wheels is captured and used.

'a major breakthrough in propulsion worldwide' and it is feasible that SKYLON could fly within a decade.[55]

When GPS was developed as a classified programme decades ago, its use in billions of cheap phone chips could not be predicted and we cannot now know much about the myriad possibilities that will come from opening up space for commerce. Two obvious applications are that it would be much easier and cheaper to: 1) monitor the earth with satellites, which brings many benefits such as improved forecasting and disaster response; and 2) protect earth from catastrophic impacts.[56] Progress with defence would overlap with commercial possibilities. Google's 'Lunar X' Prize aims to encourage private teams to deploy robots on the moon by 2015.[57] In 2012, some billionaires founded 'Planetary Resources' to pursue the long-discussed possibilities of mining asteroids and comets for valuable materials: a single platinum-rich asteroid could produce more platinum than ever mined on earth. Another application might be '*space-based solar power*' (SBSP). A 2007 Pentagon report concluded:

> 'The Sun is a giant fusion reactor ... radiating 2.3 billion times more energy [~4×10^{14} TW] than what strikes the disk of the Earth, which itself is more energy in a hour than all human civilization directly uses in a year [we use ~14 TW, the sun provides earth ~1.7×10^5 TW]... In the vicinity of Earth, every square meter of space receives [1.4] kilowatts of solar radiation, but by the time it reaches the ground, it has been reduced ... to less than an average of 250 watts per square meter... The basic idea...: place very large solar arrays into continuously and intensely sunlit Earth orbit..., collect gigawatts of electrical energy, electromagnetically beam it to Earth, and receive it on the surface for use either as baseload power ..., conversion into manufactured synthetic hydrocarbon fuels, or as low−intensity broadcast power beamed directly to consumers... A single kilometer−wide band of geosynchronous earth orbit experiences enough solar flux in one year (approximately 212 terawatt−years) to nearly equal the amount of energy contained within all known recoverable conventional oil reserves on Earth today (approximately 250 TW−yrs)... [T]there is enormous potential for energy security, economic development, improved environmental stewardship, advancement of general space faring, and overall national security...'.[58]

---

[55] Cf. Evidence to Parliament. The history of the British Government in dealing with this project since the 1980s is alarming and it would be terrible if it were to go wrong again, though the announcement of £60m funding in the 2013 Spending Review was hopeful. NanoSatisfi is developing a 10x10x10cm, 1kg, mini satellite - the CubeSat. Their goal is to lower significantly the cost of space-based scientific experiments and greatly expand the market for space commerce by shifting space from being based on i) large, monolithic, nonstandard 20-year out-of-date computer technology to ii) today's computer technology. They liken the CubeSat and its platform to the iPhone and iTunes: an iPhone is a microprocessor plus four sensors plus iTunes, and the CubeSat is a user-programmable Arduino-based microprocessor plus 25 sensors and a software interface to the web ('disruptive technology makes non-users into users'). Researchers, students, hackers *et al* will rent them. 20 years of Moore's Law has put the Cray Supercomputer in the hands of a 10 year-old with an iPad; NanoSatisfi aims to put that 10 year-old in charge of a satellite.

[56] In 1908, an asteroid ~50m wide hit Siberia, devastating an area ~150 times larger than the atomic bomb that hit Hiroshima. 'Impacts from asteroids over 1 km in diameter are expected to release over 10^5 megatons of energy with global consequences for our planet, while those with an average diameter of 100m can are expected to release over 10^2 megatons of energy potentially causing significant tsunamis and/or land destruction of a large city. It is estimated that there are between 30,000–300,000 NEO's with diameters around 100m, meaning a large number of NEO's are still undetected' (Vasile *et al*). Tracking and learning to move asteroids from their orbits is a critical objective. 'All civilizations that inhabit planetary systems must eventually deal with the asteroid threat, or they will go the way of the dinosaurs' (*Scientific American,* December 2011). It is also suggested that a formation of space-based solar power lasers could be used to move near-earth objects from threatening orbits.
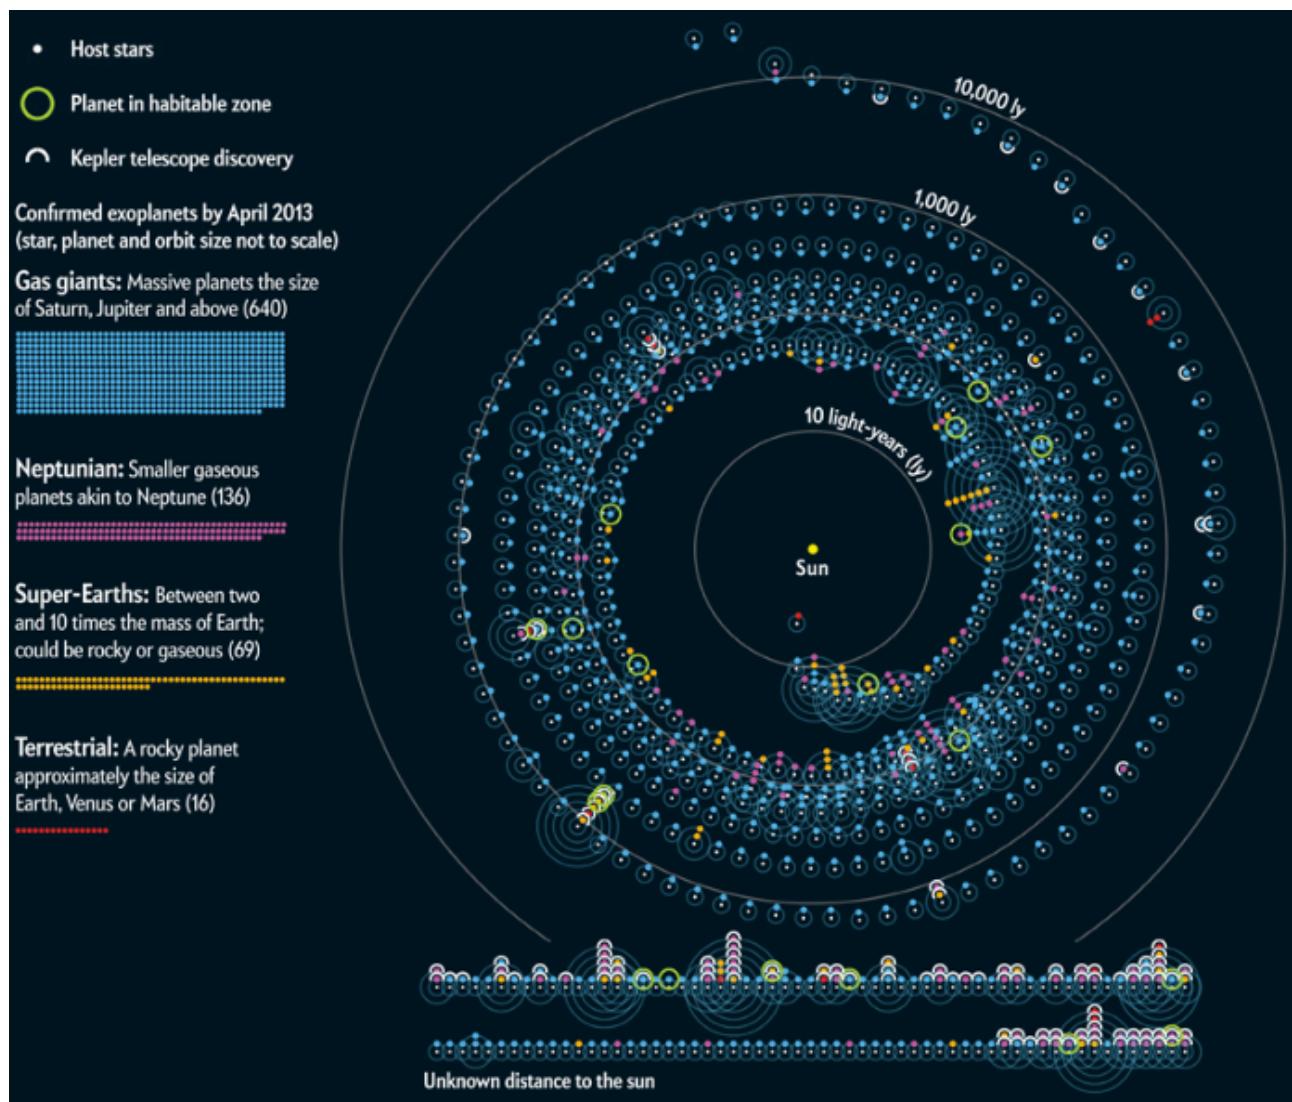
[57] Cf. '*Bound for the Moon*' (*Scientific American*, April 2012) describes the growing space industry.

[58] Japan's space agency is developing a megawatt-class SBSP demonstration. Cf. This paper on SBSP implications.

Exoplanets… In 2012, the Kepler team [announced](#) the discovery of thousands of new [exoplanets](#) including planets orbiting stars *[in the habitable zone](#)*. It is now clear that there are billions of exoplanets in our galaxy (maybe [over 1 per star on average](#) so >100 billion) and it is reasonable to infer from the recent data that there are millions of Earth-like planets orbiting habitable zones ('earth twins') in the Universe and we may soon find one. The nearest exoplanet ([found 10/2012](#)) is only 4 light-years away. In April 2013, the Kepler team [reported](#) that they had found the first plausible candidates for Earth-like planets in the habitable zone of a star. There are $\sim 10^{11}$ galaxies each containing $\sim 10^{11}$ stars so perhaps 10^{22} or more stars; this number is about a thousand times greater than the number you get by starting with 1 and doubling on each square of a 64 square chessboard ($2^{63} \approx 10^{19}$).
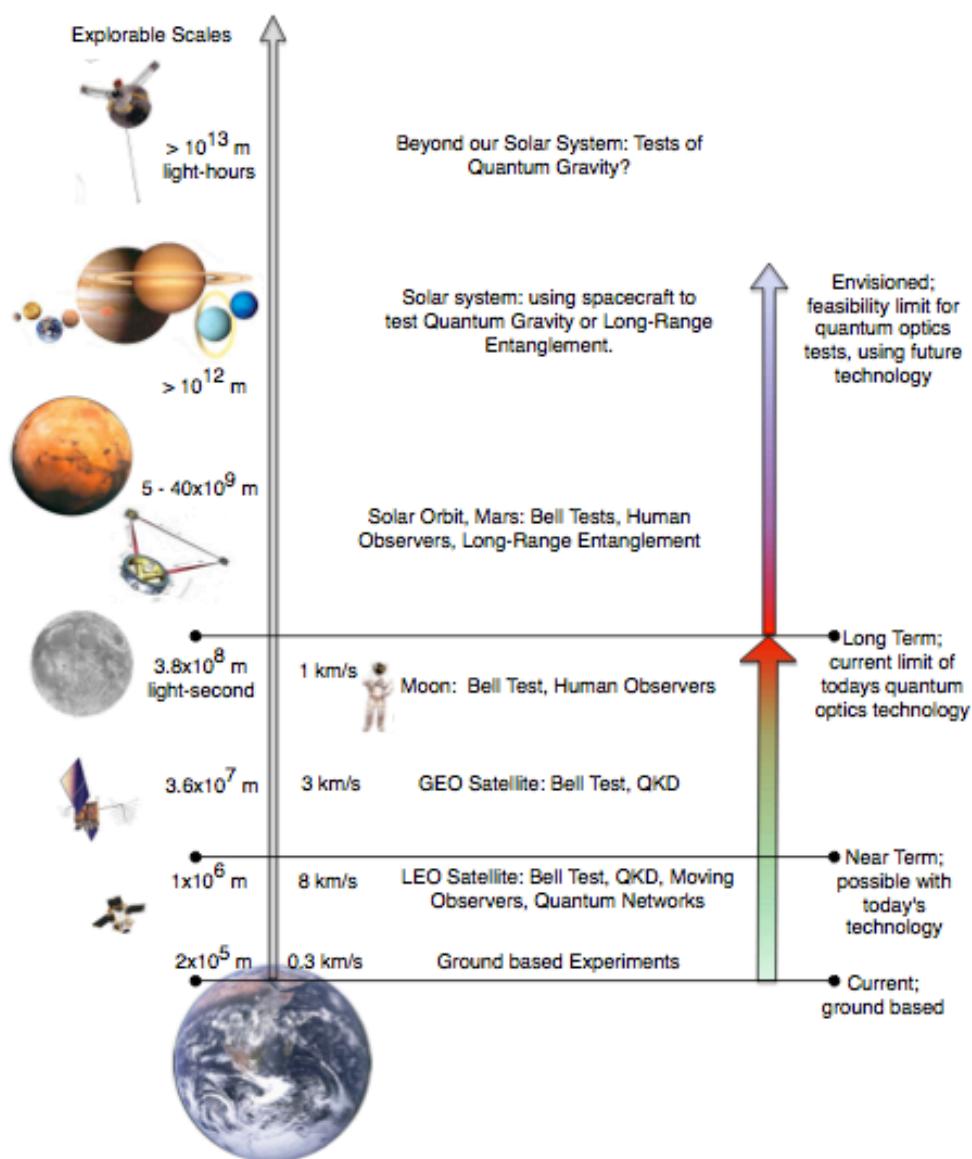
The Gaia satellite (launching 2013) is planned to map 1% of our galaxy's stars ($\sim$1 billion) and tens of thousands of exoplanets. In 2017, NASA's *Transiting Exoplanet Survey Satellite* will launch to look at half a million nearby stars (within 100 light-years). The *James Webb Space Telescope* will launch in 2018 and may be able to examine earth-like planets for signs of life ('biogenic signatures'). We should also urgently restart projects such as the [Terrestrial Planet Finder](#) and [Darwin](#) to search for such earth-like exoplanets, as detection of the chemical signatures of planets' atmospheres requires specialist technology. *Few projects have the same potential to affect mankind's psychology and behaviour.*

Graphic: Identified exoplanets as of 2013 ([Scientific American, July 2013](#))

Space and physics… It is proposed that various space experiments could help solve longstanding problems with fundamental physics including the conflicts between Relativity and Quantum Mechanics, and there is crossover between such fundamental science experiments and advances in secure communications. Quantum Mechanics was devised for atomic systems on a scale of $\sim 10^{-10}$m but has been extended downwards to $\sim 10^{-17}$-10^{-20}m in the LHC and upwards to 10^5m. Is there an upper limit to the scale on which it works, some physical boundary between the quantum and classical realms? General Relativity has been shown to apply to distant binary pulsars and to planets and satellites in our solar system. Is there a lower limit to the scale on which it works? Will a quantum theory of gravity answer both of these questions? Roger Penrose, Anders Hanssen, and others have suggested space-based experiments to probe possible boundaries between classical and quantum realms.[59] (See below Section 3 and Endnote on quantum computation.)

Figure: Overview of the distance and velocity scales of possible 'quantum information' space experiments ([Smolin et al](#)) [60]

Explorable Scales

$> 10^{13}$ m light-hours — Beyond our Solar System: Tests of Quantum Gravity?

$> 10^{12}$ m — Solar system: using spacecraft to test Quantum Gravity or Long-Range Entanglement.

Envisioned; feasibility limit for quantum optics tests, using future technology

$5 - 40 \times 10^9$ m — Solar Orbit, Mars: Bell Tests, Human Observers, Long-Range Entanglement

3.8×10^8 m light-second — 1 km/s — Moon: Bell Test, Human Observers

Long Term; current limit of todays quantum optics technology

3.6×10^7 m — 3 km/s — GEO Satellite: Bell Test, QKD

1×10^6 m — 8 km/s — LEO Satellite: Bell Test, QKD, Moving Observers, Quantum Networks

Near Term; possible with today's technology

2×10^5 m — 0.3 km/s — Ground based Experiments

Current; ground based

---

[59] It is possible that quantum gravity only manifests itself on an energy scale so high (therefore distances and times so short) as to be inaccessible - perhaps only 10^{-35}m. I.e. as far from the size of quarks (10^{-18}m) as quarks are to your car keys. (Zeilinger)

[60] QKD = Quantum Key Distribution. LEO = Low Earth Orbit. GEO = Geostationary Earth Orbit.

3. Physics and computation.

'One of the primary goals of physics is to understand the wonderful variety of nature in a unified way. The greatest advances of the past have been steps toward this goal: the unification of terrestrial and celestial mechanics by Isaac Newton in the 17th century; of optics with the theories of electricity and magnetism by James Clerk Maxwell in the 19th century; of space-time geometry and the theory of gravitation by Albert Einstein in the years 1905 to 1916; and of chemistry and atomic physics through the advent of quantum mechanics in the 1920s.' Steve Weinberg.

By about 1980, the main pillars of the 'Standard Model' of particle physics (this chart summarises the Standard Model) provided an extremely accurate (though incomplete) quantitative model of the fundamental particles (electrons, quarks etc) and three of the four fundamental forces (electromagnetism, the weak nuclear force, and the strong nuclear force) while General Relativity provided an extremely accurate quantitative model of the fourth force, gravity. However, various problems with the Standard Model remained including inconsistencies between Quantum Mechanics and Relativity.

> 'The part of the Standard Model that unites the weak and electromagnetic forces, presented in 1967–1968, is based on an exact symmetry between these forces. The W and Z particles that carry the weak nuclear forces and the photons that carry electromagnetic forces all appear in the equations of the theory as massless particles. But while photons really are massless, the W and Z [bosons] are actually quite heavy. Therefore, it was necessary to suppose that this symmetry between the electromagnetic and weak interactions is "broken" - that is, though an exact property of the equations of the theory, it is not apparent in observed particles and forces.
>
> 'The original and still the simplest theory of how the electroweak symmetry is broken, the one proposed in 1967–1968, involves four new fields that pervade the universe. A bundle of the energy of one of these fields would show up in nature as a massive, unstable, electrically neutral particle that came to be called the Higgs boson. All the properties of the Higgs boson except its mass are predicted by the 1967–1968 electroweak theory… [I]f found, it would confirm the simplest version of the electroweak theory.' (Weinberg, one of the Nobel-winners who contributed to the Standard Model, NYRB 2012.)

In 2012 the Large Hadron Collider (LHC) essentially confirmed the existence of the Higgs boson.[61]

> 'Successful as the Standard Model has been, it is clearly not the end of the story. For one thing, the masses of the quarks and leptons in this theory have so far had to be derived from experiment, rather than deduced from some fundamental principle. We have been looking at the list of these masses for decades now, feeling that we ought to understand them, but without making any sense of them. It has been as if we were trying to read an inscription in a forgotten language, like Linear A. Also, some important things are not included in the

---

[61] These two blogs, Woit's and Dorigo's, are two of the best on LHC updates.

Standard Model, such as gravitation and the dark matter that astronomers tell us makes up five sixths of the matter of the universe.' (Weinberg, op. cit.)[62]

So far the LHC has also dented the main hope of the past thirty years to solve the other problems of the Standard Model - string theory and 'supersymmetry' - against which there is a growing backlash.[63]

Two of the great 20th Century revolutions, the computer and Quantum Mechanics, have merged to produce Quantum Information Theory, Quantum Computation, and Quantum Communication. They are now attacking fundamental problems and producing new technologies. For example, despite sounding magical, '*quantum teleportation*' is not only an accepted part of quantum mechanics but has been experimentally demonstrated. It is based on the quantum phenomenon of 'entanglement': when two or more separate physical systems are 'entangled', a measurement on one of them *instantly* affects what you will find if you measure the state of the other (even if they were on opposite sides of the solar system). Einstein famously called this 'spooky action at a distance' and argued that 'no reasonable definition of reality could be expected to permit this', thus quantum mechanics must be wrong. However, all subsequent experiments have vindicated quantum mechanics (cf. Endnote).[64] Bennett showed in 1993 that although entanglement does *not* allow faster-than-light classical communication, it can be exploited for 'quantum communication'. [65]

In May 2012 the two main teams reported breakthroughs: the main Chinese team extended the range of teleportation to 100km (paper), and Zeilinger's team demonstrated entanglement over

---

[62] 'The discovery of [the Higgs] also leaves theorists with a difficult task: to understand its mass. The Higgs is the one elementary particle whose mass does not arise from the breakdown of the electroweak symmetry. As far as the underlying principles of the electroweak theory are concerned, the Higgs mass could have any value…In fact, there is something puzzling about the Higgs mass we now do observe. It is generally known as the "hierarchy problem." Since it is the Higgs mass that sets the scale for the masses of all other known elementary particles, one might guess that it should be similar to another mass that plays a fundamental role in physics, the so-called Planck mass, which is the fundamental unit of mass in the theory of gravitation... But the Planck mass is about a hundred thousand trillion times larger than the Higgs mass. So, although the Higgs particle is so heavy that a giant particle collider was needed to create it, we still have to ask, why is the Higgs mass so small?' Weinberg, 2012.

[63] Even some of string theory's supporters, such as David Gross (one of the pioneers of the Standard Model), are unhappy about some recent developments: '[T]he string theory "landscape" of 10^{500} solutions does not make sense to me. Neither does the multiverse concept or the anthropic principle, which purport to explain why our particular universe has certain physical parameters. These models presume that we are stuck, conceptually…' Though this recent paper by Nobel-winner Frank Wilczek on the Higgs defends supersymmetry and predicts the LHC will find supersymmetric particles by 2020.

[64] '*Can quantum-mechanical description of reality be considered complete?*', by Einstein, Podolsky, Rosen, *Physical Review*, 1935. Although this is one of the most cited papers in physics, Einstein was not satisfied either with his collaborators or the final state of the paper and complained to Schrödinger that, 'It did not come out as well as had originally wanted. [It] was … smothered by the formalism.' Still, he was as sure as ever, he wrote to Schrödinger, that the Copenhagen Interpretation was incomplete: 'This epistemology-soaked orgy ought to burn itself out. No doubt, however, you smile at me and think that, after all, many a young whore turns into an old praying sister, and many a young revolutionary becomes an old reactionary.' Bohr quickly responded: 'The apparent contradiction in fact discloses only an essential inadequacy of the customary viewpoint of natural philosophy for a rational account of physical phenomena of the type with which we are concerned in quantum mechanics. Indeed [the existence of quantum phenomena] entails … the necessity of a final renunciation of the classical idea of causality and a radical revision of our attitude towards the problem of physical reality.'

[65] 'The most common misconception about entanglement is that it can be used to communicate - transmit information from a sender to a receiver - perhaps even instantaneously. In fact it cannot communicate at all, except when assisted by a classical or quantum channel, neither of which communicate faster than the speed of light. So a future Internet will need wires, radio links, optical fibers, or other kinds of communications links, mostly classical, but also including a few quantum channels… I believe that the Internet will remain almost entirely classical, with quantum communication and computation being used for a few special purposes, where the unique capabilities of quantum information are needed. For other tasks, where coherent control of the quantum state is not needed, classical processing will suffice and will remain cheaper, faster, and more reliable for the foreseeable future.' Bennett.

143km (paper). Later in 2012, the Chinese team revealed that they had done the first demonstration of quantum teleportation between two macroscopic objects 150m apart (paper). In June 2013, the Chinese team announced a successful test of satellite-to-earth transmission of photons, by bouncing photons off a satellite (though some assume that classified projects have already done this given the importance of quantum cryptography). China plans to launch the 'Chinese Quantum Science Satellite' in 2016 to explore space-based quantum communication. (Zeilinger complained that European decision-making was 'so slow no decision was made' on his rival project.) The goal of such work is partly fundamental physics and partly the creation of a new satellite-based global communication network based on quantum technology. [66] Interestingly, Los Alamos recently revealed that it has been running a primitive 'quantum internet' system for a few years already.[67] Bennett recently summarised the next steps:

> 'The next and more important stage [after establishing methods for quantum communication with satellites], which depends on further progress in quantum memory and error correction, will probably be the development of a network of quantum repeaters, allowing entanglement to be generated between any two nodes in the network, and, more importantly, stockpiled and stored until needed. Aside from its benefits for cybersecurity (allowing quantum-generated cryptographic keys to be shared between any two nodes without having to trust the intermediate nodes) such a globe-spanning quantum repeater network will have important scientific applications, for example allowing coherent quantum measurements to be made on astronomical signals over intercontinental distances. Still later, one can expect full scale quantum computers to be developed and attached to the repeater network. We would then finally have achieved a capacity for fully general processing of quantum information, both locally and globally - an expensive, low-bandwidth quantum internet if you will - to be used in conjunction with the cheap high-bandwidth classical Internet when the unique capabilities of quantum information processing are needed.'

Can we solve the problems with the Standard Model, unify Relativity[68] and Quantum Mechanics into a single theory, explaining the four fundamental forces and resolving disagreements about the foundations of quantum mechanics, and provide new questions and successors to the LHC?[69] How revolutionary will Quantum Information and Computation prove, both in furthering fundamental physics and producing new technologies?

---

[66] This paper explains the importance of space for fundamental physics experiments, including problems concerning the inconsistencies of General Relativity and quantum mechanics. This paper explores an EU plan for quantum communication links to the International Space Station. This paper explores the scientific case for a moon base. An interesting detail of the Chinese satellite experiment: the satellite they claim to have used was decommissioned in 2010, suggesting they had a reason to keep this result quiet for three years. Suspicions are rife in this field because of the huge security implications of quantum communication, computation, and cryptography.

[67] In August 2013, researchers showed how to extend quantum communication to mobile devices.

[68] It has been thought that detection of gravity waves would probably have to await the LIGO (Laser Interferometer Gravitational Wave Observatory) which is being upgraded and is due to reach its design sensitivity in 2018-19. However, a recent paper suggests that gravitational waves are much stronger than previously thought and predicts detection by 2016 with 95% confidence.

[69] Part of the problem has been the growing cost of 'big science'. 'Rutherford's experimental team consisted of one postdoc and one undergraduate. Their work was supported by a grant of just £70 from the Royal Society of London. The most expensive thing used in the experiment was the sample of radium, but Rutherford did not have to pay for it - the radium was on loan' (Weinberg, op. cit). The first rudimentary 'particle accelerator' in 1932 fit on a table and split the atom successfully. The first major CERN accelerator in 1959 was 200m wide and reached 24GeV. Now, the LHC and its successors employ thousands, cost billions, and take decades from design to results.

Computation

Chart: Moore's Law 1970-2011

Microprocessor Transistor Counts 1971-2011 & Moore's Law
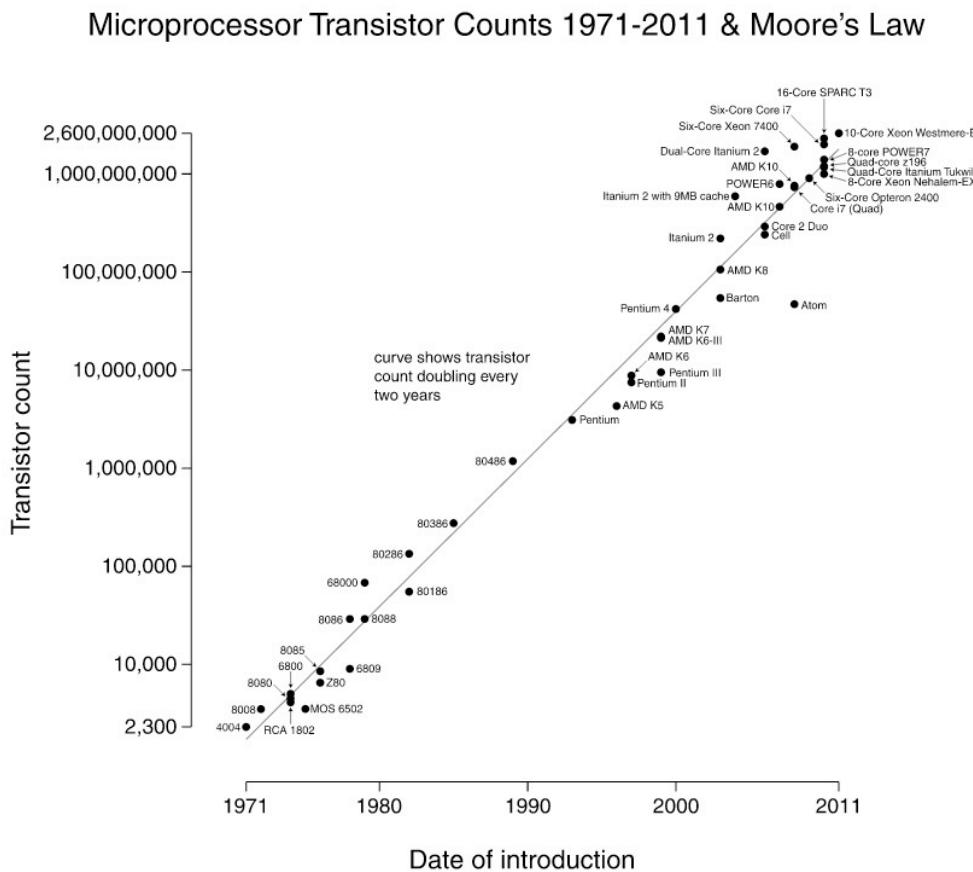
Chart: Supercomputer power - 'exaflops' (10^{18}) projected for 2018 (SUM=combined power of the top 500)
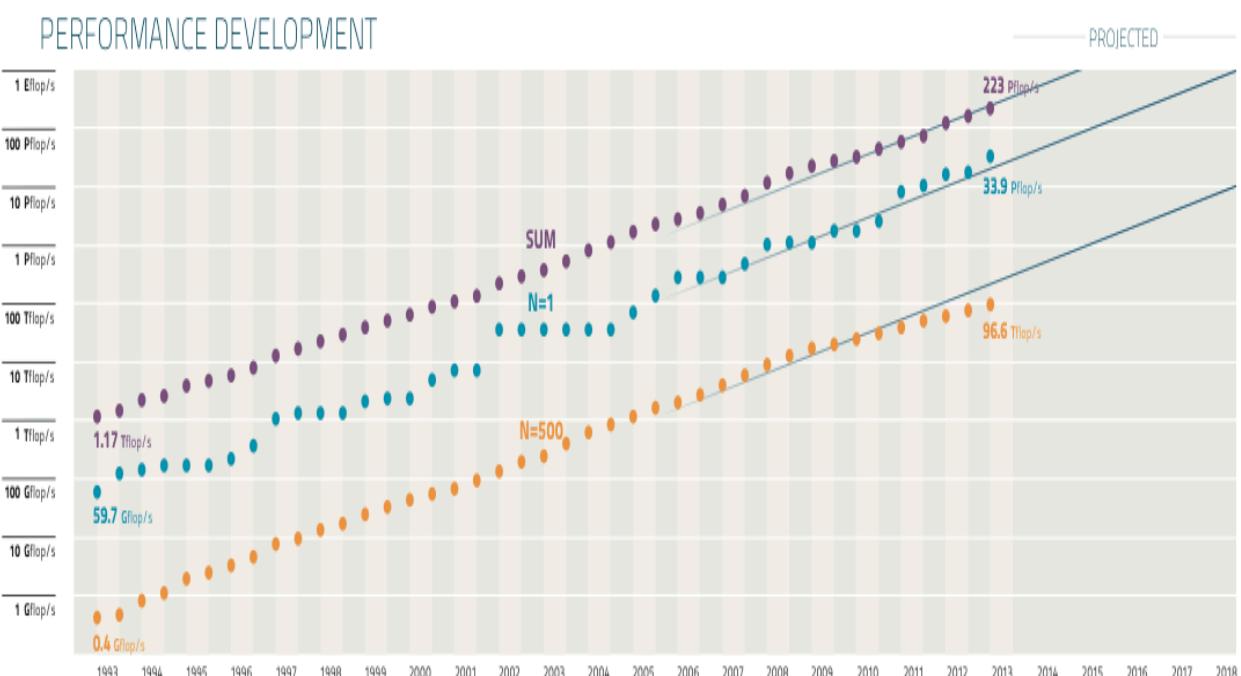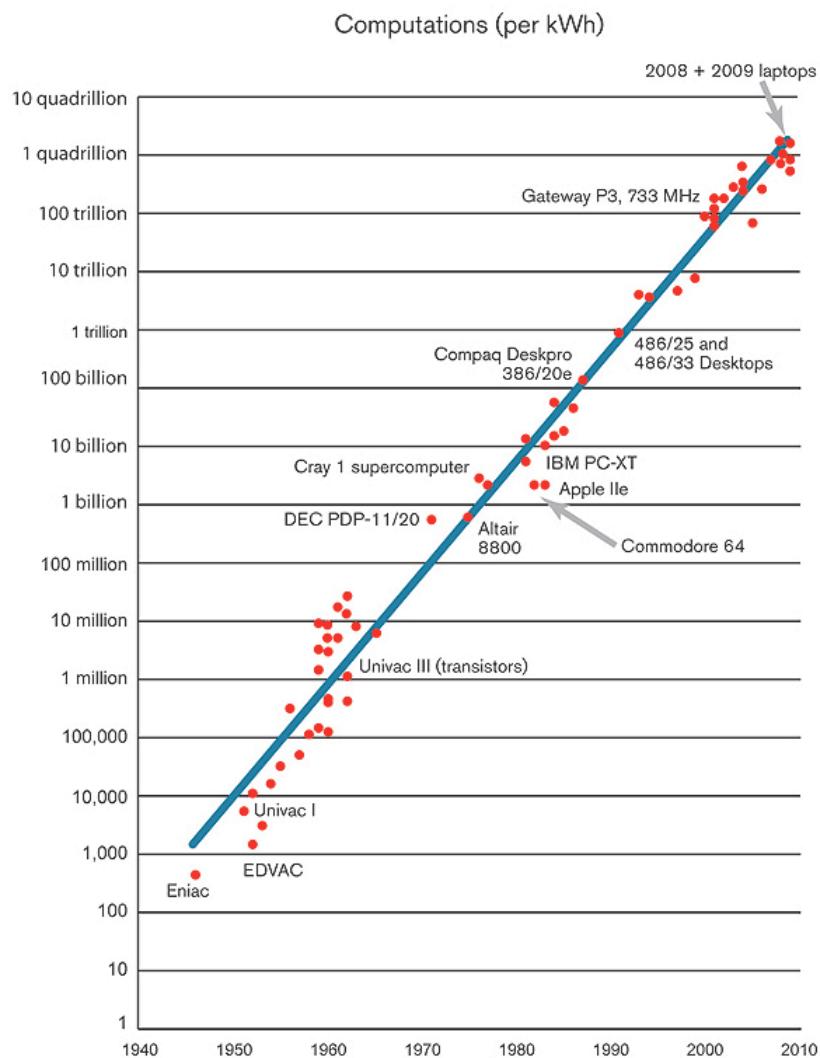
Chart: exponential spread bringing ubiquitous transistors - over 1 sextillion (10^{21}) by 2015.

Computations per kWh: The number of computations computers can carry out using the same amount of energy has been doubling every 1.5 years, so by a factor of ~100 every decade (MIT)

Desktop computers operate (2012) at ~100 gigaflops (10^{11} 'floating point operations per second'), around the level of the best *super*computer in 1993. The performance of the best supercomputer in 2013 (34 petaflops $\approx 3 \times 10^{16}$ flops) was roughly equivalent to the *sum* of the best 500 supercomputers in 2008.[70] *FoldIt* is a 'distributed supercomputer': that is, many individuals' computers are linked together *via* the internet, and it is competitive with the world's best traditional supercomputers. IBM's 'Watson', which recently beat human champions at *Jeopardy!* (see below), operates at only ~80 teraflops (8×10^{13}), hundreds of times slower than the fastest computer and about 800 times faster than a good laptop which on current trends means that a laptop would have Watson's power in about 15-20 years (10 doublings in 15-20 years = thousand-fold improvement). Theoretically, there remains massive space for improvements in computational performance (cf. Seth Lloyd's paper, *Ultimate physical limits to computation*). Moore's Law seems safe until ~2020, though by then we will be approaching the 7 nanometer scale and while Intel plans to introduce 10 nanometer technology in 2015 nobody yet knows how to push beyond 7nm..

'Exascale computation' (10^{18} or 1 quintillion flops) from ~2018-20 is expected to bring breakthroughs in: fundamental physics; designing and manufacturing materials, including at the nanoscale; understanding complex biochemical cycles, protein engineering, genotype>phenotype interactions; more accurate modelling of various turbulent systems such as weather. Will exascale computing be able to simulate the brain? The brain consists of ~80 billion (8×10^{10}) neurons and ~100 trillion (10^{14}) synapses. Operating on ~15-20 watts it performs ~10^{17} floating point computations per second. Section 5 describes some of the projects underway to prepare for exascale computing to simulate the brain.

Increasing performance is connected to improving computation per unit of energy, miniaturisation, heat loss, and fundamental architecture. The 2012 top performing supercomputer (now #2) required 8MW to deliver its ~20 petaflops.[71] However, there has been a doubling in computations per unit of energy every 1.5 years since World War II so an improvement in performance by a factor of ~100 every decade. This has already brought sensors that can scavenge energy from tiny stray background energy flows such as TV signals or ambient heat. If it continues for another two decades, then it will mean today's most powerful supercomputers operating not with the power needed to run a town, as now, but with the power of a toaster, and it would mean the computational power of today's laptops requiring 10,000 times less power than now, with all that implies for further miniaturisation. Smartphones are roughly as powerful as the supercomputers of 25 years ago and if trends continue a computer as powerful as today's iPhone could be the size of a cell in a few decades.

However, continued miniaturisation involves dealing with heat. Transistors have been packed onto ever denser 2-dimensional chips and this means that a lot of heat is emitted by the resistance of the wires transporting information (much more than is emitted by the logic switching of the transistors). As the amount of information processed per watt grows, so does the amount of heat generated. Increasingly efficient cooling systems have therefore been necessary but nothing that exists now could sufficiently cool an exascale computer. One solution is already happening - a switch from 2D to 3D architecture with transistors stacked in three dimensions. 'If you reduce the linear dimension by a factor of ten, you save that much in wire-related energy, and your information arrives almost ten times faster' (Michel, IBM).

---

[70] In 1964, the CDC 6600 operated at 3×10^6 flops (3 megaflops). In 1976, the Cray 1 operated at ~1.4×10^8 (136 megaflops). In 1985, the Cray 2 operated at 1.9×10^9 flops (1.9 gigaflops). In 1996, the SR2201 operated at 6×10^{11} flops (600 gigaflops). In 2002, the NEC Earth Simulator operated at 3.5×10^{13} flops (35 teraflops). In 2007, the IBM Blue Gene operated at 6×10^{14} flops (600 teraflops). In 2010, the Tianhe-1A operated at 2.5×10^{15} flops (2.5 petaflops).

[71] The 2013 #1, China's Tianhe-2 requires ~18 MW to deliver its 34 petaflops.

This switch to 3D architecture is connected to the work on replacing the 'von Neumann architecture' that all computers have been based on since von Neumann's pioneering work in 1945 - a switch to an architecture based on the brain's efficient operation. The brain accounts for just 2% of the body's volume but 20% of its total energy demand. On average neural tissue consumes about ten times more power per unit volume than other human tissues. However, relative to computers the brain is extremely efficient.If an exascale machine were to require just ~20 megawatts, that is still a million times more than the ~20 Watts required by the brain.

'In a computer, as much as 96% of the machine's volume is used to transport heat, 1% is used for communication (transporting information) and just one-millionth of one percent is used for transistors and other logic devices. By contrast, the brain uses only 10% of its volume for energy supply and thermal transport, 70% for communication and 20% for computation. Moreover, the brain's memory and computational modules are positioned close together, so that data stored long ago can be recalled in an instant. In computers, by contrast, the two elements are usually separate. "Computers will continue to be poor at fast recall unless architectures become more memory-centric" [Michel].'

'[A brain-like] hierarchy is already implicit in some proposed 3D designs: stacks of individual microprocessor chips (on which the transistors themselves could be wired in a branching network) are stacked into towers and interconnected on circuit boards, and these, in turn, are stacked together, enabling vertical communication between them. The result is a kind of 'orderly fractal' structure, a regular subdivision of space that looks the same at every scale.

'Michel estimates that 3D packaging could, in principle, reduce computer volume by a factor of 1,000, and power consumption by a factor of 100, compared to current 2D architectures. But the introduction of brain-like, 'bionic' packaging structures, he says, could cut power needs by another factor of 30 or so, and volumes by another factor of 1,000. The heat output would also drop: 1-petaflop computers, which are now large enough to occupy a small warehouse, could be shrunk to a volume of 10 litres.

'If computer engineers aspire to the awesome heights of zetaflop computing (10^{21} flops), a brain-like structure will be necessary: with today's architectures, such a device would be larger than Mount Everest and consume more power than the current total global demand. Only with a method such as bionic packaging does zetaflop computing seem remotely feasible. Michel and his colleagues believe that such innovations should enable computers to reach the efficiency - if not necessarily the capability - of the human brain by around 2060.' (Ball, '*Feeling the heat*', *Nature* 2012. See also below Section 5.)

While most media attention focuses on Moore's Law and hardware, progress in some software has also been dramatic.

'The algorithms that we use today for speech recognition, for natural language translation, for chess playing, for logistics planning, have evolved remarkably in the past decade. It's difficult to quantify the improvement, though, because it is as much in the realm of quality as of execution time. In the field of numerical algorithms, however, the improvement can be quantified. Here is just one example... Grötschel, an expert in optimization, observes that a benchmark production planning model solved using linear programming would have taken 82 years to solve in 1988, using the computers and the linear programming algorithms of the day. Fifteen years later – in 2003 – this same model could be solved in roughly 1 minute, an improvement by a factor of roughly 43 million. Of this, a factor of roughly 1,000 was due to increased processor speed, whereas a factor of roughly 43,000 was due to improvements in

algorithms! Grötschel also cites an algorithmic improvement of roughly 30,000 for mixed integer programming between 1991 and 2008.'[72] (Cf. Endnote re Linear Programming.)

The fundamental architecture for almost all contemporary computers comes from Turing's 1936 logical vision (cf. Endnote), in which the computer 'head' (processor) is separate from the 'tape' (information), and von Neumann's 1945 engineering architecture ('*The First Draft of a Report on the EDVAC*'). This architecture keeps the hardware and software separate: the hardware is physical and software is virtual. Both von Neumann and Turing, having created one approach and seen it take root, knew that it had limitations and spent their last few years before their untimely deaths thinking about different models that may be possible. For example, Von Neumann's '*Theory of Self-Reproducing Automata*' launched the field of 'cellular automata' and first proved the concept of a self-replicating machine.

As the original Turing/von Neumann architecture spread, another tradition was developed by Landauer, Bennett and others as they worked on fundamental issues of physics, information, computation, and thermodynamics (including nagging problems with Maxwell's Demon). Wheeler coined the phrase 'It from Bit' and Landauer described how 'information is physical': that is, *information is registered and processed across all physical systems and all physical systems register and process information.*[73] The relationship between thermodynamics, information, and computation was described by Charlie Bennett: 'Computers may be thought of as engines for transforming free energy into waste heat and mathematical work'.[74] Theorists and engineers, particularly at MIT (e.g. Seth Lloyd and Neil Gershenfeld) and IBM (e.g. Charlie Bennett), have built on this tradition. There are four related developments: *quantum computation* (dealt with mainly in the Endnote); *digital fabrication* in the sense of conventional computers used to programme personal fabrication machines; pushing digital fabrication to the atomic scale with *atomically precise manufacturing*; and the elimination of the the distinction between hardware and software in such developments as *programmable matter*.

In a famous 1981 talk (written up as '*Simulating Physics With Computers*'), Richard Feynman discussed some of the problems with modern physics, including the difficulty of simulating quantum mechanics on a classical (i.e. Turing/von Neumann) computer and speculated on the possibility of 'quantum computation'.

> 'During the 1960s and 1970s, Richard Feynman was involved in attempts to use classical digital computers to evaluate the consequences of quantum field theory. He observed that quantum mechanics was hard to program on a classical digital computer. The reason for this difficulty was straightforward: Quantum mechanics possesses a variety of strange and counterintuitive features, and features that are hard for human beings to comprehend are also hard for classical computers to represent at the level of individual classical bits. Consider that a relatively small quantum system consisting of a collection of 300 electron spins "lives" in $2^{300} \approx 10^{90}$ dimensional space. As a result, merely writing down the quantum state of the spins in a classical form as a sequence of bits would require a computer the size of the

---

[72] Report by President's Council of Advisors on Science and Technology. The success of the 'open source software' community, based on the principle that 'Given enough eyeballs all bugs are shallow', is stunning. Cf. 2 blogs (1,2) by a leading computer scientist on the importance of properly framing the question an algorithm is intended to solve.

[73] '*The Physical Nature of Information*', Landauer 1996.

[74] '*The Thermodynamics of Computation*', Bennett 1981. In '*Demons, Engines and the Second Law*' (*Scientific American*, 1987), Charles Bennett explained for the non-specialist how the problems raised by Maxwell's Demon, which had rumbled from the 1870s through the quantum and computer revolutions, had finally been tamed by Landauer: saving the second law of thermodynamics required arguments from quantum mechanics and computer science.

universe, and to compute the evolution of that state in time would require a computer much larger than that.

'Feynman noted that if one has access to a quantum-mechanical device for representing the state of the spins and for transforming that state, rather than a classical device, then the computation of the time evolution of such a system can be much more economical. Consider a collection of 300 two-level quantum systems, or qubits [quantum bit = 'qubit'], one for each electron spin. Suppose that one sets up or programs the interactions between those qubits to mimic the dynamics of the collection of spins. The resulting device, which Feynman called a universal quantum simulator, will then behave as a quantum analog computer, whose dynamics form an analog of the spin dynamics.' ([Lloyd, 2008]).

The merging of quantum mechanics and computer science leads to a fundamental insight: *our physical universe computes quantum mechanically and this is much more efficient than a classical computer*.[75] This insight provokes a natural question: can we exploit this to build quantum computers that would be much more efficient than classical computers? This is touched upon in an Endnote on Quantum Computation and the history of this thinking is explained by Gershenfeld in this fascinating [talk].

We have developed digital communication (~1945) and digital computation (~1955), and now we are beginning another revolution - *digital fabrication*. Mainframes were expensive with limited markets, needed skilled operators in special rooms and performed repetitive industrial operations. Now, the machines that make computers (fabricators) are expensive with limited markets, need skilled operators in special rooms and perform repetitive industrial operations. Just as mainframes evolved into minicomputers then PCs and became accessible to individuals in a mass market, so fabricators have evolved into an intermediate stage and are becoming accessible to individuals in a mass market (personal fabricators). Individuals gained control of *bits* during the 1980s and 1990s with the PC revolution. They are now taking control of *atoms* with the PF revolution. (Cf. Gershenfeld's *Fab*)[76]

Gershenfeld and colleagues started a course at MIT in 1998 called '*How to make (almost) anything*.' It was overwhelmed by demand. Students would teach others and produce their own teaching materials. They then developed 'fab labs' as 'proto-personal fabricators' to push the revolution forward. They have spread across the world and are used in places such as Afghanistan to make vital tools locally. Each originally cost $20,000 and Gershenfeld's Fablabs are what he describes as the 'minicomputer' phase of the PF revolution.

The PC was driven by apps such as VisiCalc (the first spreadsheet was a huge boost for Apple in 1979) and Lotus 1-2-3 (which boosted IBM in 1983). The PF revolution is being driven by *the falling cost of precise control of time and space*: sign cutters, tabletop mills, laser cutters and waterjet cutters (subtraction), and 3D printers (e.g. shoot a laser into a bath of polymer that hardens where the laser hits the surface, building up a part in layers) which print objects such as insulators,

---

[75] '*Computational Capacity of the Universe*', Lloyd 2002.

[76] At MIT, 'the Whirlwind' was developed in the late 1940s and demonstrated in 1951. It introduced display screens to give an instant output. The USAF supported a project to connect it to an industrial milling machine to help make aircraft. From this there developed a new kind of programming language for 'computer-aided manufacturing' (CAM) with numerically controlled machines. The first was Automatically Programmed Tools which was available on IBM's 704 in 1958. The TX-2 at MIT, using transistors instead of vacuum tubes, had a 'light pen' that 'allowed an operator to draw directly on the display screen'. In 1960, one of Shannon's students used the combination of the TX-2 and light pen to create the 'Sketchpad' program that allowed a designer to 'sketch shapes which the computer would then turn into precise geometrical figures'. It was the first 'computer-aided design' (CAD) program. (Cf. Gershenfeld's '*Fab*')

conductors, and semiconductors (addition). This allows precise control of hardware that previously was very expensive.

> 'Laboratory research... has shown how to print semiconductors for logic, inks for displays, three-dimensional mechanical structures, motors, sensors, and actuators. We're approaching being able to make one machine that can make any machine… [A] little table-top milling machine can measure its position down to microns … so you can fabricate the structures of modern technology such as circuit boards for components in advanced packages. And a little 50-cent microcontroller can resolve time down below a microsecond, which is faster that just about anything you might want to measure in the macroscopic world. Together these capabilities can be used to emulate the functionality of what will eventually be integrated into a personal fabricator...

> 'Powders and plastics that conduct electricity can be used to print wires, there are printable semiconductors that can be used to deposit logic circuits, motors can be made with magnetic materials, combinations of chemicals can store energy, and so forth. Printable inks containing each of these types of materials have been developed… Integrating all of them into a printer is the most promising route toward making one machine that can make anything. The joke about a student graduating when his or her thesis can walk out of the printer follows from that goal - the student must be able to print the text, the structural elements, the actuators for motion, the control systems for logic, and a supply of energy.' (Gershenfeld)

The frontier of Computer Aided Design (CAD) is changing the computer/fabricator interface, such as with 'data gloves' (the man who did the special effects in Minority Report is making a real version). 'A mathematical specification, a graphical rendering, a tabletop prototype, and a full-size structure can now contain exactly the same information just represented in different forms… The intersection of 3D scanning, modelling, and printing blurs the boundaries between artist and engineer, architect and builder, designer and developer, bringing together not just what they do but how they think' (Gershenfeld).

The modern economy was based on mass production rather than personalised production and modern education separated maths, logic, science, and computers from design and production. Both are changing. Personal fabricators are shrinking in price dramatically.[77] Personal fabrication together with the 'open source' movement, and tools such as Alibaba and Kickstarter that enable large, decentralised collaborations and *democratised mass production* (i.e. personal fabrication scaled up), promises economic,[78] educational (see Section 6), and military[79] revolutions. Coase described the purpose of the firm as reduced transaction costs (the people you need are brought together) but as Bill Joy said, 'No matter who you are, most of the smartest people work for someone else'; this new model makes collaboration with the best people, wherever they are, much easier (Anderson). We can also reunite making things with the training of the intellect in our education: 'Instead of trying to interest kids in science as received knowledge, it's possible to equip them to

---

[77] A MakerBot is <$1,000. Recently, such technology cost ~$25-50,000.

[78] The movie *Flash of Genius* tells a 1960s story. A professor invents a windscreen wiper and makes a prototype. He does a deal with Ford and builds a factory. However, Ford steal his idea and drop the deal. The professor is bankrupt. Now, the professor could email the prototype design to www.alibaba.com, it would be drop-shipped to his door, and he could tweak the prototype before scaling up production *via* an outsourced manufacturer and ship it direct to customers. (Chris Anderson, lecture).

[79] 3D printers enable big improvement in price/performance of DIY drone technology (this is possible by exploiting regulatory loopholes otherwise forbidden by ITAR). People can now print automatic weapons at home.

do science, giving them both the knowledge and the tools to discover it' (Gershenfeld, cf. Section 6).

Programmable matter… Now, the most advanced factories use digital computers to manipulate analog materials. Even the most sophisticated chip labs use a process 'relatively unchanged from what artisans have done for millennia: layers of materials are deposited, patterned, and baked. The materials themselves are inert.' Airplane manufacture is similar; 'the actual fabrication process happens by a rotating piece of metal whacking away at a stationary piece of metal. The intelligence is external to the tools' (Gershenfeld). In the future we will, Gershenfeld argues, remove the distinction between hardware and software and computing will become 'a bulk raw material' that you buy by the kilogram. 'Digitising fabrication doesn't mean a computer is connected to a tool - it means the computer *is* a tool. It doesn't mean a program describes a thing - it means a program *is* a thing. The information is in the materials, not external to the system' (Gershenfeld). This also means '*programmable matter*' - *materials whose properties can be programmed*, where the distinction between object and programme dissolves (like with DNA).[80]

> 'Every day, scientists and engineers design new devices to solve a current problem. Each device has a unique function and thus has a unique form. The geometry of a cup is designed to hold liquid and is therefore different from that of a knife which is meant to cut. Even if both are made of the same material (e.g., metal, ceramic, or plastic), neither can perform both tasks. Is this redundancy in material, yet limitation in tasks entirely necessary? Is it possible to create a programmable material that can reshape for multiple tasks? *Programmable matter* is a material whose properties can be programmed to achieve specific shapes or stiffnesses upon command… This concept requires constituent elements to interact and rearrange intelligently in order to meet the goal.' ('*Programmable matter by folding*', Hawkes *et al* (2010).)

Digital fabrication will make intelligence *internal* to the fabrication process. Just as Shannon showed that you can produce reliable communication from unreliable signals, and allowing digital communication, and von Neumann showed that you can have reliable computation with unreliable components, allowing digital computation, so 'digital fabrication will enable perfect macroscopic objects to be made out of imperfect microscopic components'. This will require 'the development of an associated design theory that can specify what an enormously complex machine should do without specifying in detail how it does it' (Gershenfeld).

A 2013 paper from Gershenfeld shows a breakthrough combining 'digital materials' and digital fabrication. Instead of forming large composite carbon fibre structures in conventional ways, Gershenfeld has shown that they can be made from many identical, tiny, interlocking parts that can be mass-produced. This produces much stronger and lighter structures that also reduce the cost and increase the design flexibility. When conventional structures fail (often at joints), they tend to do so catastrophically. The new structures, exploiting new geometries, behave like an elastic solid that is massively redundant and is therefore less likely to fail catastrophically.

> '[These materials] can be considered as a "digital" material: a discrete set of parts is linked with a discrete set of relative positions and orientations. These attributes allow an assembler to place them using only local information rather than global positioning, and allow placement errors to be detected and corrected. These attributes are familiar in nano-assembly, and are used here on macroscopic scales. These materials combine the size and strength of composites with the low density of cellular materials and the convenience of

---

[80] C.f. '*Reconfigurable Asynchronous Logic Automata*', Gershenfeld et al, 2010).

additive manufacturing. The need for custom tooling is eliminated, because parts can be incrementally added to a structure. Their construction, modification, repair, and re-use can all employ the same reversible linking process. Heterogeneous elements can be incorporated in structures with functions determined by their relative placement. Exact assembly of discrete cellular composites offers new properties and performance not available with the analog alternatives of continuously depositing or re- moving material.'

'Internet of things', ubiquitous computing… Another aspect of the transition to digital fabrication is the connection of the physical world to the internet. Gershenfeld *et al* developed the idea of '[the internet of things](#)' to break down the division between the physical world and the internet. It is very wasteful for it to cost a thousand dollars to add a dumb sensor to a building and to have central computers control the physical environment rather than the users. We need to bring the principles of the internet to the physical world: 1) 'it's incrementally extensible, you add to it at the edges'; 2) 'the applications live at the edges so what the network does is defined by what you connect to it' (Gershenfeld).

Moore's Law has spawned an exponential spread of transistors (chart above). There were ~5 quintillion (5×10^{18}) transistors by 2005. Between 2005-2010 this grew about 15 times to ~75 quintillion - about 10 billion for every person on earth, dwarfing the number of grains of rice harvested. Intel predicts (2012) that by 2015 there will be 1,200 quintillion (1.2 sextillion = 1.2×10^{21}) transistors. This is making sensor networks ubiquitous, stretching from the ocean floor to the atmosphere. Large numbers of high density sensors (light, temperature, humidity, vibration and strain) + energy-scavenging ability (e.g. solar or thermoelectric device, wireless power) + GPS embedded + ultra-sensitive accelerometer (sensor tracks the motion of a tiny internal movable platform relative to the rest of the chip and measures changes in acceleration) + internet connection = disruption: for example, infrastructure monitoring, health monitoring, chemical / biological hazard monitoring, 'smart buildings', intrusion / illicit movement detection.

New materials… Manipulating materials at the nanoscale can bring huge advantages. Carbon fibre is a fifth of the mass of steel and is just as strong and its price is down to a thirtieth of its 1970 price, hence it was used for the space shuttle and is now used in bicycles and golf clubs. Carbon is not rare - it is the fourth most abundant material in the universe - so if we can make further innovations, then we have the potential for immense breakthroughs. We already have two candidates - *graphene* and *carbon nanotubes*. A one-atom thick array of carbon atoms in the form of graphene is virtually see-through, 300 times stronger than steel and 1,000 times better conductor than silicon (graphene is the best room temperature conductor we know). Carbon nanotubes are similarly superior to steel and normal conductors; if one shoots a hole in them they self-repair in femtoseconds. The smallest are about 0.7 nm so inside is the quantum world and we are engineering at the limits of our understanding ([Bill Joy](#)). Chip manufacture is already at the scale of <50 nanometres and people are experimenting with graphene transistors.

Ultimately this unification of physics and computation in 'digital fabrication' may enable 'atomically precise fabrication' and 'nanotechnology' first described in another famous [talk](#) by [Feynman](#), *There's Plenty of Room at the Bottom (1959)*. The US National Research Council recently [assessed](#) the feasibility of atomically precise manufacturing and the sort of roadmaps needed, and one of the pioneer's [blog](#) tracks developments.

*

4. Biological engineering.

'If they lived just one month amid the misery of the developing world, as I have for fifty years, they'd be crying out for tractors and fertilizer and irrigation canals and be outraged that fashionable elitists back home were trying to deny them these things.' Norman Borlaug, leader of the Green revolution, Nobel Peace Prize winner.

'In 1953, when the structure of DNA was determined, there were 53 kilobytes of high-speed electronic storage on planet earth. Two entirely separate forms of code were set on a collision course… Everything that human beings are doing to make it easier to operate computer networks is at the same time, but for different reasons, making it easier for computer networks to operate human beings... Darwinian evolution ... may be a victim of its own success, unable to keep up with non-Darwinian processes that it has spawned.' George Dyson.

Chart: US market penetration of GM crops (Carlson)
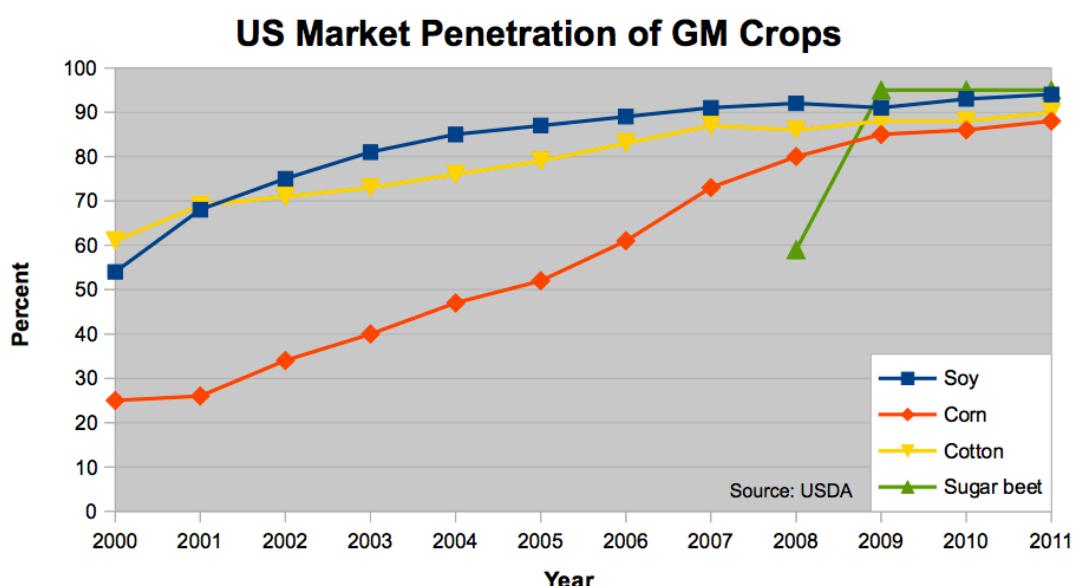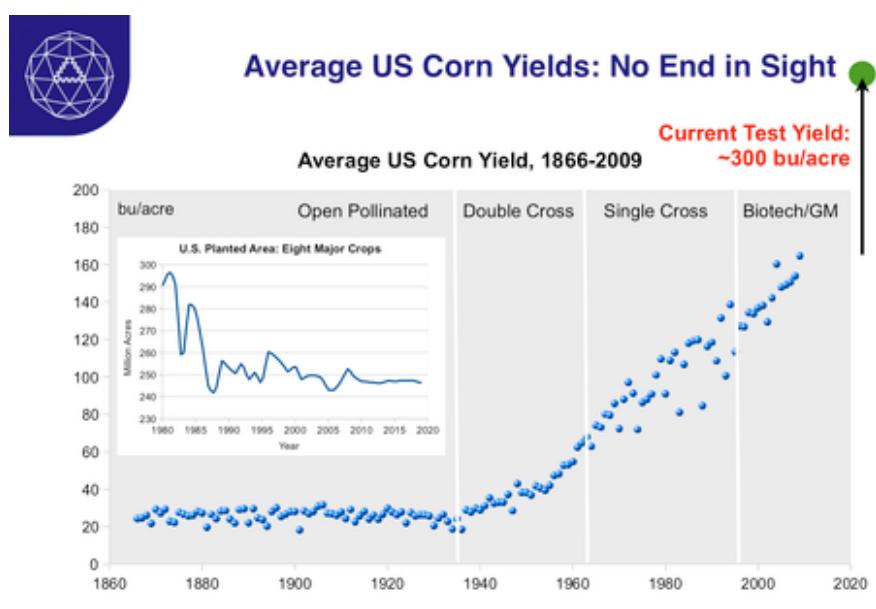
Chart: Average US Corn Yields (Carlson)
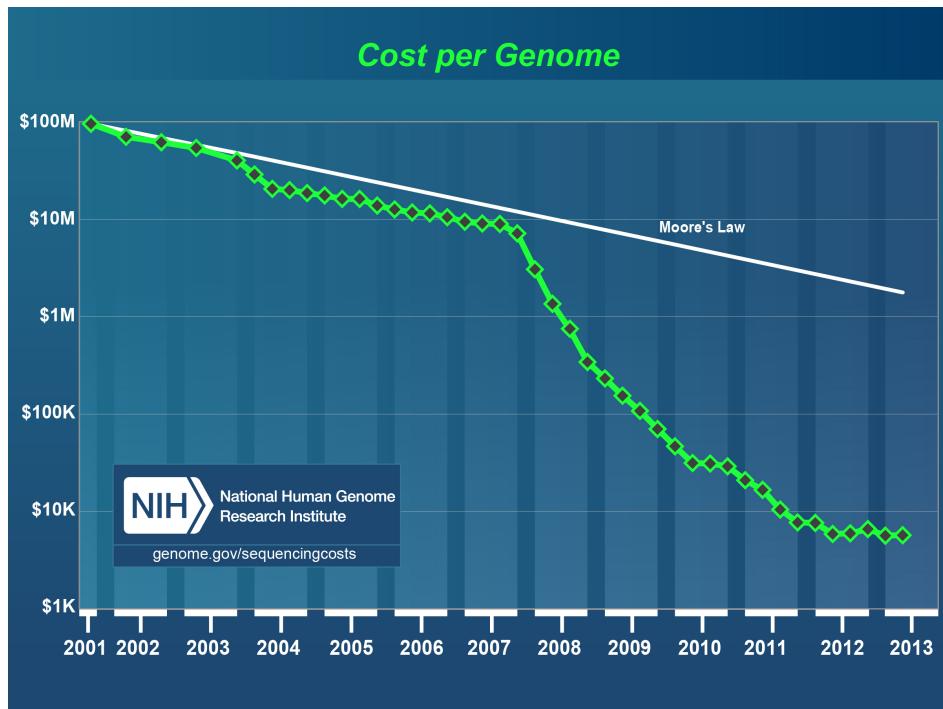
Chart: ~10,000- fold improvement in price/performance of DNA sequencing vs ~16-fold improvement in Moore's Law

Faster than Moore's Law acceleration in the speed of genome sequencing has enabled progress from a) a decade and billions of dollars to sequence the first human genome (1990-2003) to b) months and millions, to c) days and thousands now, and, soon, d) hours and hundreds of dollars per human genome.[81]

Genetically modified crops could help feed the world. 'Golden rice' (with genes modified so it grows vitamin A) has been developed for years and is now being grown in trials; it could significantly cut millions of deaths and blindness caused by vitamin A deficiency. Scientists have also modified Africa's most common staple crop, cassava, into 'BioCassava', a variant with increased levels of vitamin A, iron, and protein. Some plants use a type of photosynthesis that evolved more recently and is more efficient ('C4 photosynthesis') and the Gates Foundation is studying how these more efficient genes could be transferred to rice and wheat. This could bring rice and wheat that produce 'one and a half times the yield per acre, that require less water per calorie, that need less fertilizer per calorie, and that are more resistant to drought.' Many similar projects are underway. Crops modified genetically to resist pests or pesticides (like RoundUp) lower the environmental impact of agriculture.[82] Overall, GMOs could allow more people to be fed healthily while lowering our environmental impact: significantly lower water and pesticide use, less energy needed, less pressure on remaining forests and so on. Every major scientific study has concluded that there is

---

[81] A sudden acceleration from early 2008 saw us break the $1m per genome barrier in 2008, the $100,000 barrier in 2009, and the $10,000 barrier in 2011. As of July 2013, it costs <$6,000. BGI will soon (2014?) be sequencing >1,000 genomes per day at <$1,000 per genome. In 2012, Nanopore announced a huge breakthrough (disposable USB drive sequencers) that might deliver the $100 genome even faster (cf. the $100 genome by JASON for DOD).

[82] Organic food is not a solution. While it reduces some environmental effects (e.g pesticides), organic food has significantly lower yields so it requires more land.

no evidence of ill health from eating GMOs and the occasional 'evidence' quoted by Greenpeace *et al* is misinterpreted.[83] (Naam, 2013.)

Biomedicine is advancing. We are developing new drug delivery mechanisms, genetic engineering, and biometric analysis for personalised medicine. Sequencing the genes of cancer cells ('whole genome sequencing') has had dramatic results recently. '*Induced pluripotent stem cells*' (iPSCs) allow reprogramming of cells into different types which is contributing to the industry for growing replacement organs: ears, noses and rat hearts have already been made; in July 2013 it was reported for the first time 'the generation of vascularized and functional human liver from human iPSCs by transplantation of liver buds created in vitro… [This is] the first report demonstrating the generation of a functional human organ from pluripotent stem cells... [T]his proof-of-concept demonstration of organ-bud transplantation provides a promising new approach to study regenerative medicine' (*Nature, 2013*); in 2013 it was also shown 'conclusively that [embryonic stem cells] can provide a source of photoreceptors for retinal cell transplantation' which may soon bring a cure for 'irreversible' blindness caused by loss of photoreceptors (*Nature, 2013*). '*RNA interference*' allows turning off gene expression.[84] '*Optogenetics*' allows switching neurons on/off with lasers and is transforming neuroscience. In 2010, there was a demonstration of the first transplantation of a *synthetic genome* into a recipient cell: an organism had been created with DNA constructed in a computer. Nanotechnology and robotics are merging in medicine and will bring *greater precision,* instead of reliance on traditional methods of cutting (surgery), burning (radiotherapy), and poisoning (chemotherapy).[85]

After the molecular biology revolution (post-Crick & Watson) and the genomics revolution (post-2000) comes 'convergence', the merging of technologies, disciplines or devices that creates new possibilities. It involves the integration of fields - particularly subfields of *physical sciences, engineering, and life sciences*. Physical science has already been transformed by *IT, materials, imaging, nanotechnology, optics, quantum physics* together with advances in *computing, modelling, and simulation*. Convergence takes the tools traditional to physics and engineering and applies them to life sciences creating new fields like *bioinformatics*, *computational biology*, and *synthetic biology*.[86]

---

[83] Various groups such as Greenpeace and Friends of the Earth have successfully campaigned against GMOs. Rich activists who shop in London and New York have encouraged African countries to reject GM foods even during famines and praised those regimes that listened to them (e.g Zambia).

[84] In July 2013, it was announced (*Translating dosage compensation to trisomy 21*) that for the first time genetic engineering has been used to silence the effects of the extra chromosome that causes Down's syndrome.

[85] In 1966, Fantastic Voyage publicised the idea of tiny sensors inside us (remade as Innerspace 1987). By 1987 we were building protypes of pill-size robots. In 2000 patients began swallowing the first pill cameras, moved *via* external magnets. Now pill robots can power themselves and take samples. Modular robots could self-assemble, work, disassemble and act with *precision.* 'Every pharmaceutical and biotech company makes chemical or biological molecules that target [biochemical] structures… No one is designing medicines to interact with the electrical signals that are the other language of biology. The sciences that underpin bioelectronics are proceeding at an amazing pace… The challenge is to integrate the work - in [BCIs], materials science, nanotechnology, micro-power generation - to provide therapeutic benefit.' Head of R&D at GSK announcing new investments, 2012.

[86] E.g. *Computational biology* integrates: computer science, physics, engineering, molecular and genetic biology. Simultaneously, life science models (e.g. viral self-assembly) are influencing physical sciences and engineering. A major breakthrough in the modelling of a cell at the molecular level was announced July 2012, 'a whole-cell computational model of the life cycle of ... *Mycoplasma genitalium* [a tiny bacterium that operates with ~430 genes] that includes all of its molecular components and their interactions.'

How soon will we have precise enough understanding and tools to understand the biological foundations of human nature and cognition, repair biological failure, re-engineer species, or 'enhance' humans?[87]

'*Missing heritability*'… Twin and adoption studies have showed that many fundamental human qualities are at least 50% heritable (i.e at least half the variance between individuals is caused by genes): e.g. 'g' or 'general cognitive ability' (~60-80% heritable for adults) and some personality characteristics. After the success of the Human Genome Project over a decade ago, there was much hope of finding 'the gene for X'. However, while some rare conditions and disorders (including learning disorders) are caused by a mutation in a single gene, most common traits and disorders (including learning disorders) are caused by *many genes with small effects*, making them hard to find.[88] This became known as 'the missing heritability problem': we know that genes are responsible but we cannot find many of the specific genes. The combination of *whole-genome sequencing* and '*Genome-Wide Association Studies*' (*GWAS*) is likely to allow us to make progress in finding this 'missing heritability'.

First, the fall in cost means that *whole-genome sequencing* will soon be viable for millions. Only ~2% of the genome (~20,000 genes) codes for proteins. For years the rest of the genome was referred to as 'junk DNA'. However, we now know that much of this 98% is not 'junk' but has various little understood functions, such as switching on/off the protein coding genes. A decade-long study reported (2012) that more than 50 percent of the genome is biochemically active. Although many rare single-gene disorders involve mutations in coding genes, noncoding genes may contribute to the heritability of complex traits and common disorders that are the object of much of the behavioural sciences. The former director of the Human Genome Project, Francis Collins, said (2010) that 'I am almost certain … that whole-genome sequencing will become part of newborn screening in the next few years.' If it does become standard for all new babies, then vast amounts of genome data will become available for no extra cost.

Second, (GWAS), enabled by 'microarrays' (chips the size of a postage stamp) which can now cheaply genotype over a million DNA variants across the genome, allow us to search the entire genome for correlations, among unrelated individuals, between an allele and a trait (i.e. individuals with a particular allele differ on the trait from those with different alleles). The first GWAS was published in 2005 and *Science* hailed it as 'Breakthrough of the Year' in 2007. GWAS have already confirmed the findings of the twin and adoption studies of various characteristics, and begun to identify genes associated with weight, height, cognitive ability and other things though only a small part of the 'missing heritability' has so far been found (and some studies have not been replicated

---

[87] Cf. This study by JASON for America's DOD. According to a survey by *Nature*, scientists are already using drugs like Modafinil to gain a competitive edge. We have recently sequenced the Neanderthal genome, *contra* many predictions, and Dr. George Church (Harvard Medical School) said that a Neanderthal could be brought to life with present technology for about $30 million and that if offered funding he would consider doing it. Dr. Klein (Stanford) commented, 'Are you going to put them in Harvard or in a zoo?'

[88] 'The essence of Fisher's insight was that genes could work hereditarily as Mendel's experiments indicated, but if a trait is influenced by several genes it would be normally distributed as a quantitative trait. The converse of this insight is important in the present context: If a disorder is influenced by many genes - and this is the conclusion that emerges from GWA research - its genetic liability is likely to be normally distributed. Thus, in terms of genetic liability, common disorders are quantitative traits. In other words, genes that are found to be associated with disorders in case-control studies are predicted to be correlated with the entire range of variation throughout the normal distribution. Stated more provocatively, this means that from a genetic perspective there are no common disorders, just the extremes of quantitative traits' (Plomin, '*Child Development and Molecular Genetics*', 2013).

suggesting that some findings have been false positives).[89]

We are learning more about how specific genes affect specific behaviour (Weber *et al*, Nature 2013) and prediction of phenotypes from genome information is *already* happening with inbred cattle. The continued fall in sequencing costs to <$10,000 then <$1,000 per genome and the ever-larger GWAS samples, growing from 10^3 to 10^6 people, give grounds for confidence that over the next decade we will begin to find the missing heritability for personality, common disorders, and various aspects of cognition.

One such project is the BGI project (working with Robert Plomin) to identify *the genes responsible for 'general cognitive ability' or 'g'*. BGI has the largest capacity to sequence genomes in the world and is sequencing the genomes of thousands of people with IQ scores of +3 standard deviations (1:1,000) to identify the genes responsible. One of the team is physicist Steve Hsu who explains the state of the art in intelligence and IQ research in a talk here. Even more radically, Hsu thinks that once the genes are identified, then engineering higher intelligence might become feasible (Plomin is much more cautious about this). He quotes a famous comment by the Nobel-winning physicist and mathematician Wigner about von Neumann:

> 'I have known a great many intelligent people in my life. I knew Planck, von Laue and Heisenberg. Paul Dirac was my brother in law; Leo Szilard and Edward Teller have been among my closest friends; and Albert Einstein was a good friend, too. But none of them had a mind as quick and acute as Jansci von Neumann. I have often remarked this in the presence of those men and no one ever disputed me… Perhaps the consciousness of animals is more shadowy than ours and perhaps their perceptions are always dreamlike. On the opposite side, whenever I talked with the sharpest intellect whom I have known - with von Neumann - I always had the impression that only he was fully awake, that I was halfway in a dream.'

Hsu says: 'I'm doing my best to increase the number of future humans who will be "fully awake". My current estimate is that one or two hundred common mutations (affecting only small subset of the thousands of loci that influence intelligence) are what separate an ordinary person from a vN [von Neumann].'[90] In May 2013, a GWAS study reported (*Science*) the first 'hits' for genetic variants associated with educational achievement.

---

[89] 'GWA research might eventually identify more than half of the heritability. This criterion fits with another criterion for GWA success: to exceed the prediction from family data. First-degree relatives are on average 50% similar for additive genetic effects so that identifying half of the heritability should exceed the prediction that is possible from first-degree relatives. However, DNA predictions can be more valuable than family risk estimates for three reasons. First, DNA predictions are specific to individuals within a family in contrast to family risk estimates, which are the same for all members of a family. Second, predictions based on DNA are limited to genetics whereas predictions from family risk can include nurture as well as nature. Third, DNA sequence variation does not change during development whereas family risk estimates - for example, using parents' characteristics to predict children's risks - can be complicated by developmental change' (Plomin, '*Child Development and Molecular Genetics*', 2013).

[90] 'The most extensive selection experiment, at least the one that has continued for the longest time, is the selection for oil and protein content in maize (Dudley 2007). These experiments began near the end of the nineteenth century and still continue; there are now more than 100 generations of selection. Remarkably, selection for high oil content and similarly, but less strikingly, selection for high protein, continue to make progress. There seems to be no diminishing of selectable variance in the population. The effect of selection is enormous: the difference in oil content between the high and low selected strains is some 32 times the original standard deviation' (Crow). Selection can produce changes faster than people realise. The Russian scientist, Belyaev, began an experiment to simulate the evolutionary process of domestication. He took wild silver foxes and selectively bred them on the basis solely of *tameability* (only a small fraction were allowed to breed). It has run for over four decades. The result is that the original population has been transformed from wild animals (*via* 40-45 generations of breeding) to tame animals with various physical changes (including body size, markings, and reproductive cycles) and biochemical changes (e.g altered serotonin levels). Cf. Endnote on IQ.

Debate over such issues is clouded by the lack of understanding of many social scientists about the basic science and statistical methods and this lack of understanding seeps into the media (e.g. Gladwell's '*Outliers*' repeated [many false ideas](#) about 'g' and 'IQ' that have circulated) partly because those working at the cutting edge in genetics are understandably reluctant to involve themselves in contentious debates and partly because, for understandable reasons given our history, people are wary of discussing the importance of genes in explaining human behaviour (cf. Section 6 on education and Endnote on genetics and intelligence).

*

5. Mind and Machine.

Many types of robot are already deployed from the operating theatre to the battlefield.[91] DARPA's 2002 Grand Challenge for autonomous vehicles transformed the field (Google's driverless cars are already on America's roads); their 'robot mule' can now track a soldier and they are getting faster.[92] In April 2012 they announced a new Robotics Grand Challenge for robots operating in a disaster environment. This project offers prizes for competing teams to design and deploy 'supervised autonomous' humanoid robots to perform basic tasks, including using human tools, to solve problems in a 'disaster' environment. Winners of the first stage will be given an ATLAS robot (unveiled June 2013) - a state-of-the-art humanoid robot - to programme for the final Grand Challenge in 2014. If it is as successful as the 2002 Grand Challenge, robotics will make another huge advance.

> 'This challenge is going to test supervised autonomy in perception and decision-making, mounted and dismounted mobility, dexterity, strength and endurance in an environment designed for human use but degraded due to a disaster. Adaptability is also essential because we don't know where the next disaster will strike. The key to successfully completing this challenge requires adaptable robots with the ability to use available human tools, from hand tools to vehicles.'[93]

In 1945, the entire power of the most educated country on earth, Germany, was insufficient to get a single bomber across the Atlantic; recently, a 77 year-old blind man built his own drone that crossed the Atlantic. On 9/11, the US had a handful of drones; by 2013 the US had deployed over 20,000 and is now training more unmanned system operators than fighter and bomber pilots combined. However, this exponential spread of drones is at a relatively early stage compared to the car revolution after 1900; we are roughly at the stage when people started to think about basic rules like 'drive on the left' and 'speed limits' for Ford's Model-T. The first car bomb was in 1905 and the first plane hijacking in 1931 - both very early in the technology's development. Given drones are developing faster than cars, we have a lot of hard problems to solve quickly (Singer).

Drones are piggybacking the smartphone industry's exploitation of better and cheaper chips, sensors (gyroscopes, magnetometers, pressure sensors, accelerometers, GPS), batteries, and optics. Autopilots piggyback the progress of the iPhone's ARM chip. The main difference between the cheap civilian drone and the Global Hawk is that the former fly at hundreds of feet for ~1 hour and the latter fly at 60,000 feet for days. The aircraft industry's shift to unmanned drones, free of the design constraints of human passengers, may bring large efficiency gains. Drones are already

---

[91] The military has been the source of much progress in computers and networks since Turing and von Neumann. Possibly the first major project to build a computerised communication and decision-making system was the USAF's SAGE which von Neumann helped develop. Three of the main figures in the history of the internet and modern personal computers are Licklider (papers), Engelbart ('*Augmented Human Intelligence*', 1960), and Baran who worked at ARPA and RAND in the 1960s. In 1968, Engelbart presented his team's ideas at a conference in one of the most famous presentations in the history of computer science. He demonstrated word processing, video chat, real-time collaborative document editing, hypertext links and the mouse. This was revolutionary in a world of teletypes and punch cards and showed a path to the personal computer revolution. One announcement was ARPA's idea of a new 'experimental network' which became the internet. Cf. Endnote.

[92] Cf. Singer's '*Wired for War*'. It took about 500 hoplites to cover a football field area; Civil War soldiers covered that area with 20; World War I soldiers covered that area with two soldiers; a single World War II soldier covered five fields; in Iraq in 2008, 'the ratio of personnel to territory was roughly 780 football fields per one US soldier' (Singer).

[93] The May 2013 *Scientific American* reported recent breakthroughs with human-robot collaboration on manufacturing projects. Instead of a big robot tucked away behind safety gates, they can now work within arm's reach using 'augmented reality' (projecting images and text directly onto objects). One experiment saw a computer science graduate and a robot significantly outperform experts in building a frame for an armoured vehicle. These new robots can be taught a new process by watching a human do it and applying machine learning algorithms.

used by film companies, farmers (which [Chris Anderson thinks](#) will be the biggest growth area in the next five years), and scientists for aerial maps. Drones have already been used by Taiwanese jewel thieves, Mexican drug cartels, and terrorist groups such as Hamas and Hezbollah. Small drones can infiltrate an area and hack cellphone calls and SMS messages without the knowledge or help of the service provider. *Matternet* plans to use transportation drones in countries without infrastructure and *Pirate Bay* is building swarms of file-sharing drones. Perhaps we will soon see political campaigns use drones to monitor opponents or protesters use them for counter-surveillance of police, who are building their own drone swarms.[94] Autonomous machines, such as automated sniper systems in Korea and Israel, have *already* been deployed to kill people without reference to human decisions. 'Swarming' has a long history (e.g. Themistocles at Salamis and the Mongols' 'Falling Stars') and remains a cutting-edge approach to conflict (cf. Arquilla's '[*Swarming*](#)'): it has already been applied to cyberwar ('botnet swarms') and now it is being applied to autonomous lethal robots (including HYDRA, a global undersea swarm of unmanned platforms). Researchers are even [controlling drones](#) *via* brain-computer-interfaces.[95] (Cf. [Singer, 2013](#).)

'Machine intelligence' (or 'automated reasoning') has made dramatic breakthroughs, from [chess](#)[96] to language [interfaces and translation](#), to [understanding natural language](#), and in 2011 IBM's '[Watson](#)', a 80 teraflop supercomputer, defeated human champions in the quiz show *Jeopardy!* and is now being put to work on medicine and finance. Automated reasoning 'bots' make decisions about your face, CV, mortgage, assets,[97] dating, [fMRI scans](#), shopping, health insurance records, your plane or car, and communication and surveillance networks. They can take and give orders *via* voice recognition software. They are increasingly built into the physical environment around you including clothing ('haptics'). They are used to [destroy things](#) in the physical world[98] and for [espionage](#) and '[cyberwar](#)'. Given that intelligent adversaries aim to implement Sun Tzu's dictum 'to win without fighting is the highest form of war', cyberwar and other advanced technologies offer obvious possibilities for 'winning' by destroying an adversary's desire, or ability, to oppose - *without conventional conflict and even without any obvious evidence.*[99]

---

[94] Counter-drone systems in the white, grey, and black worlds will involve regulation, disruption, and destruction.

[95] Cf. [Chris Anderson on the drone revolution](#). The ability to use networks to control objects and humans, combined with other technologies (e.g. encrypted 'darknets') strengthens individuals/groups versus states, cf. the brilliant sci-fi novel [*Daemon*](#), and this [talk](#) by its author Suarez. The ease of blocking electromagnetic signals may spur the rapid spread of lethal autonomous drone swarms using visual intelligence systems to identify, track and destroy targets. Suarez's '*Kill Decision*' focuses on anonymous drone swarms and some of what he describes is [already happening](#) in Obama's White House. 'In its vision of flying robots that are mass-produced like cheap toys and smart enough to think for themselves, Suarez's fiction is closer to reality than most people think' (Chris Anderson). IARPA is also [developing](#) silent drones. Cf. This [talk by Vijay Kumar](#) on robot 'swarms', using lessons from the decentralised coordination of autonomous ants using local information. DARPA is working with iRobot on swarms of modular robots and Harvard's microrobotics lab is working on swarms of insect-sized robots ('[RoboBees](#)'). [Luc Steels](#) is developing robots that evolve their own language. A [2013 paper by Levesque](#) is an interesting critique of many AI projects.

[96] By 1994, a computer beat the world checkers champion; in 2007 Shaeffer proved that every game of checkers must end in a draw with perfect play and released an unbeatable program. In 1985, Kasparov played 32 games simultaneously against the best chess computers in the world and won all 32 games. In 1997, IBM's Deep Blue beat Kasparov in a close match. In 2006, Deep Fritz trounced Kramnik. Progress is rapid in Go and poker.

[97] In 1995, Dave Cliff (University of Bristol) devised one of the first autonomous adaptive algorithmic trading systems that outperformed humans in experimental financial markets. Hedge funds like Renaissance now do this routinely.

[98] The use of *Stuxnet* to destroy Iranian nuclear infrastructure is the most obvious example but the trend is spreading. On 15 August 2012, the data on three-quarters of Saudi Arabia's ARAMCO (oil company) was destroyed in what was 'probably [the most destructive [cyber]attack that the private sector](#) has seen' (then Defense Secretary Leon Panetta).

[99] DARPA is [looking](#) for ways to automate solutions to ancient problems of spies stealing documents as well as automate cyberwar itself (e.g. '[Plan X](#)' to 'create revolutionary technologies for understanding, planning, and managing cyberwarfare in real-time, large-scale, and dynamic network environments ... and support development of fundamental strategies and tactics needed to dominate the cyber battlespace'). Cf. JASON's study '[*Science of cyber security*](#)' (2010).

There have been many approaches to 'machine intelligence' including the growing use of probabilistic models based on vast data sets (see Section 6 on Nielsen).[100] Another approach is 'genetic algorithms' that mimic biological evolution to *evolve* answers to very complex problems where exhaustive searches are not feasible.[101] They are even being applied to the 'discovery of physical laws directly from experimentally captured data', and to primitive '[robot scientists](#)'.

In 2009, Schmidt *et al* reported successfully developing an algorithm that can discover mathematical relationships in data (*[Distilling Free-Form Natural Laws from Experimental Data,](#)* Science, 2009).

> 'For centuries, scientists have attempted to identify and document analytical laws that underlie physical phenomena in nature. Despite the prevalence of computing power, the process of finding natural laws and their corresponding equations has resisted automation. A key challenge to finding analytic relations automatically is defining algorithmically what makes a correlation in observed data important and insightful. We propose a principle for the identification of nontriviality. We demonstrated this approach by automatically searching motion-tracking data captured from various physical systems, ranging from simple harmonic oscillators to chaotic double-pendula. Without any prior knowledge about physics, kinematics, or geometry, the algorithm discovered Hamiltonians, Lagrangians, and other laws of geometric and momentum conservation. The discovery rate accelerated as laws found for simpler systems were used to bootstrap explanations for more complex systems, gradually uncovering the "alphabet" used to describe those systems.

> 'Mathematical symmetries and invariants underlie nearly all physical laws in nature, suggesting that the search for many natural laws is inseparably a search for conserved quantities and invariant equations…

> '[A genetic algorithm is used.] Initial expressions are formed by randomly combining mathematical building blocks such as algebraic operators {+, −, ÷, ×}, analytical functions (for example, sine and cosine), constants, and state variables. New equations are formed by recombining previous equations and probabilistically varying their subexpressions. The algorithm retains equations that model the experimental data better than others and abandons unpromising solutions. After equations reach a desired level of accuracy, the algorithm terminates, returning a set of equations that are most likely to correspond to the intrinsic mechanisms underlying the observed system…

> 'We have demonstrated the discovery of physical laws, from scratch, directly from experimentally captured data with the use of a computational search. We used the presented approach to detect nonlinear energy conservation laws, Newtonian force laws, geometric invariants, and system manifolds in various synthetic and physically implemented systems without prior knowledge about physics, kinematics, or geometry. The concise analytical expressions that we found are amenable to human interpretation and help to reveal the physics underlying the observed phenomenon. Many applications exist for this approach, in

---

[100] DARPA has [launched](#) *The Probabilistic Programming for Advanced Machine Learning* program 'to advance machine learning by using probabilistic programming to: (1) dramatically increase the number of people who can successfully build machine learning applications, (2) make machine learning experts radically more effective, and (3) enable new applications that are impossible to conceive of using today's technology.' [Its HACMS program](#) aims to create a new architecture for software writing and verification.

[101] How does a GA work? Define the problem; create a population of diverse software agents; let the agents compete in a digital 'fitness landscape'; successful agents breed using sexual reproduction and crossover, so differential replication improves overall fitness; mutation throws some kinks in to stop a system settling into local equilibria; iterate many generations; see where the population converges in the fitness landscape; examine the evolved 'genotype' to see why it's successful; maybe throw in some random noise (eg. climate changes) or 'virtual predators'.

fields ranging from systems biology to cosmology, where theoretical gaps exist despite abundance in data.

'Might this process diminish the role of future scientists? Quite the contrary: Scientists may use processes such as this to help focus on interesting phenomena more rapidly and to interpret their meaning.'[102]

This has already yielded breakthroughs: for example, in 2013, this method was used to analyse a huge astronomy data set and 'derive an analytic expression for photometric redshift' (*The first analytical expression to estimate photometric redshifts suggested by a machine*, Krone-Martins *et al*, 2013).

Similarly, King *et al* built a 'robot scientist' that automates the hypothetico-deductive method and the recording of experiments. It makes hypotheses, devises experiments to test them, runs the experiments, interprets the results, and repeats the cycle.[103]

'The basis of science is the hypothetico-deductive method and the recording of experiments in sufficient detail to enable reproducibility. We report the development of Robot Scientist "Adam," which advances the automation of both. Adam has autonomously generated functional genomics hypotheses about the yeast *Saccharomyces cerevisiae* and experimentally tested these hypotheses by using laboratory automation. We have confirmed Adam's conclusions through manual experiments. To describe Adam's research, we have developed an ontology and logical language.'

They argue that science has developed using natural language but 'scientific knowledge is best expressed in formal logical languages' because 'only formal languages provide sufficient semantic clarity to ensure reproducibility and the free exchange of scientific knowledge.' The combination of computers, the semantic web, computerised ontologies and formal languages makes 'the concept of a robot scientist' possible.

'This is a physically implemented laboratory automation system that exploits techniques from the field of artificial intelligence to execute cycles of scientific experimentation. A robot scientist automatically originates hypotheses to explain observations, devises experiments to test these hypotheses, physically runs the experiments by using laboratory robotics, interprets the results, and then repeats the cycle...

'Adam's hardware is fully automated such that it only requires a technician to periodically add laboratory consumables and to remove waste. It is designed to automate the highthroughput execution of individually designed microbial batch growth experiments in microtiter plates. Adam measures growth curves (phenotypes) of selected microbial strains (genotypes) growing in defined media (environments). Growth of cell cultures can be easily measured in high-throughput, and growth curves are sensitive to changes in genotype and environment...

'To set up Adam for this application required (i) a comprehensive logical model encoding knowledge of S. cerevisiae metabolism … expressed in the logic programming language Prolog; (ii) a general bioinformatic database of genes and proteins involved in metabolism; (iii)

---

[102] Three further 2012 papers by Lipson ([1](#),[2](#), [3](#)) on use of algorithms to discover mathematical laws.

[103] '*The Automation of Science*', *Science*, April 2009. The pioneer experiment in this field was Stanford's DENDRAL in the 1960s to prepare for a Mars mission.

software to abduce[104] hypotheses about the genes encoding the orphan enzymes, done by using a combination of standard bioinformatic software and databases; (iv) software to deduce experiments that test the observational consequences of hypotheses (based on the model); (v) software to plan and design the experiments, which are based on the use of deletion mutants and the addition of selected metabolites to a defined growth medium; (vi) laboratory automation software to physically execute the experimental plan and to record the data and metadata in a relational database; (vii) software to analyze the data and metadata (generate growth curves and extract parameters); and (viii) software to relate the analyzed data to the hypotheses; for example, statistical methods are required to decide on significance. Once this infrastructure is in place, no human intellectual intervention is necessary to execute cycles of simple hypothesis-led experimentation.'

They report various modest but valuable discoveries made by their primitive robot scientist in functional genomics and they think that improved software could yield much deeper breakthroughs. In '*Towards Robot Scientists for autonomous scientific discovery*' (Sparkes, King *et al*, 2010) they explain how their successor to Adam, called 'Eve', will be used to automate drug screening. Other groups are pursuing similar ideas: Lipson (Cornell) is using automated experiments to improve mobile robotics; Rabitz (Princeton) is trying to automate the use of femtosecond lasers to make or break chemical bonds.[105]

Mathematicians and computer scientists are collaborating on problems concerning the use of computers to find new ideas, write proofs, and verify proofs. For example, in 1998 Thomas Hales used a computer to solve the 400-year-old Kepler Conjecture concerning the densest way to stack spheres. However, the *Annals of Mathematics* could not verify the correctness of the 50,000 lines of computer code. Hales is therefore creating new tools for formal computer verification of proofs (the Flyspeck Project) based on set theory (cf. Hales' essay '*Mathematics in the Age of the Turing Machine*'). Vladimir Voevodsky is working on a new programme called 'Coq' to verify proofs. Fields Medallist Tim Gowers also recently announced that he is developing a programme to create a mathematical theorem-*creating* computer programme. In time, it is speculated that programmes to verify proofs will step-by-step bootstrap themselves to greater conceptual depth and give us access to automated high level mathematical reasoning. (Cf. '*In Computers We Trust?*')

> 'The aim is to develop technologies until formal verification of theorems becomes routine at the level of ... Perelman's proof of the Poincaré conjecture... So far at least, no computer proof has defied human understanding. But eventually, we will have to content ourselves with

---

[104] What is abduction? An example of *deduction* is the syllogism: all swans are white, Daisy is a swan, therefore Daisy is white. An example of *abduction* is: all swans are white, Daisy is white, therefore Daisy is a swan. This is not logically valid; it is a stab in the dark, a guess that can be used as the basis for a hypothesis.

[105] Cf. *Rise of the Robot Scientists*, King (*Scientific American*, January 2011). Nobel winner Frank Wilczek thinks that by 2100 the leading physicist will be a robot.

fables that approximate the content of a computer proof in terms that humans can comprehend' (Hales).[106]

Neuroscience, cognitive science, behavioural genetics, and evolutionary biology have developed our understanding of the mind. They and other disciplines have combined with Moore's Law and scanning technology to provide increasingly accurate maps and quantitative models of the brain[107] (which Obama promised federal support for in 2013) and rapidly improving brain-computer interfaces.[108] We can look inside our minds with tools that our minds have given us and watch the brain watching itself. These developments have undermined the basis for Descartes' *Ghost in the Machine*, Locke's *Blank Slate*, and Rousseau's *Noble Savage (pace* the current, sometimes misguided, backlash).[109]

The brain consists of ~80 billion (8×10^{10}) neurons and ~100 trillion (10^{14}) synapses. Operating on ~20 watts it performs ~10^{17} floating point computations per second. As well as the brain-computer interfaces already underway, various projects plan to map the brain completely (e.g the Connectome Project), simulate the brain (e.g Markram's project), and build new computer architectures modelled on the brain and performing similarly to the brain.

---

[106] The issue of computer verification (or creation) of proofs raises two further fundamental issues: first, the issue of *trust and hacking* (*quis custodiet ipsos custodes*); second, so-called '*soft errors*'. Computers are *not* absolutely infallible. Leaving aside hardware glitches and software bugs, which can be fixed, there are 'soft errors' such as physical interference. E.g. Energetic neutrons strike the nucleus of an atom in a computer and throw off an alpha particle that hits a memory circuit and changes a value (this problem creates three times more errors in Denver than Boston due to elevation). 'Soft errors and susceptibility to hacking have come to be more than a nuisance to me. They alter my philosophical views of the foundations of mathematics. I am a computational formalist – a formalist who admits physical limits to the reliability of any verification process, whether by hand or machine. These limits taint even the simplest theorems, such as our ability to verify that 1 + 1 = 2 is a consequence of a set of axioms. One rogue alpha particle brings all my schemes of perfection to nought. The rigor of mathematics and the reliability of technology are mutually dependent; math to provide ever more accurate models of science, and technology to provide ever more reliable execution of mathematical proofs' (Hales, op. cit).

[107] Cf. Interview and book of one of the leaders of the Connectome project to map the brain precisely. A full electron microscope scan of the brain would produce ~1 zettabyte (10^{21}), which $\approx$ imaging the entire world at a resolution of one byte per m^2 for a month. 'All movie DVDs released to date comprise about a PetaByte and the text for all books ever published is "only" 30 terrabytes.' (JASON). 'I imagine that the speed of mapping connectomes will double every year or two. If so, then it will become possible to map an entire human connectome within a few decades… Using new methods of light microscopy, neurophysiologists are now able to image the signals of hundreds or even thousands of individual neurons at the same time… Imagine knowing the activity and connectivity of all the neurons in a small chunk of brain. This capability is finally within reach, and is bound to revolutionize neuroscience.' (Seung). Paul Allen is funding a project to map a mouse's brain. A recent paper (Boyden *et al*, 2013) summarises the challenges.
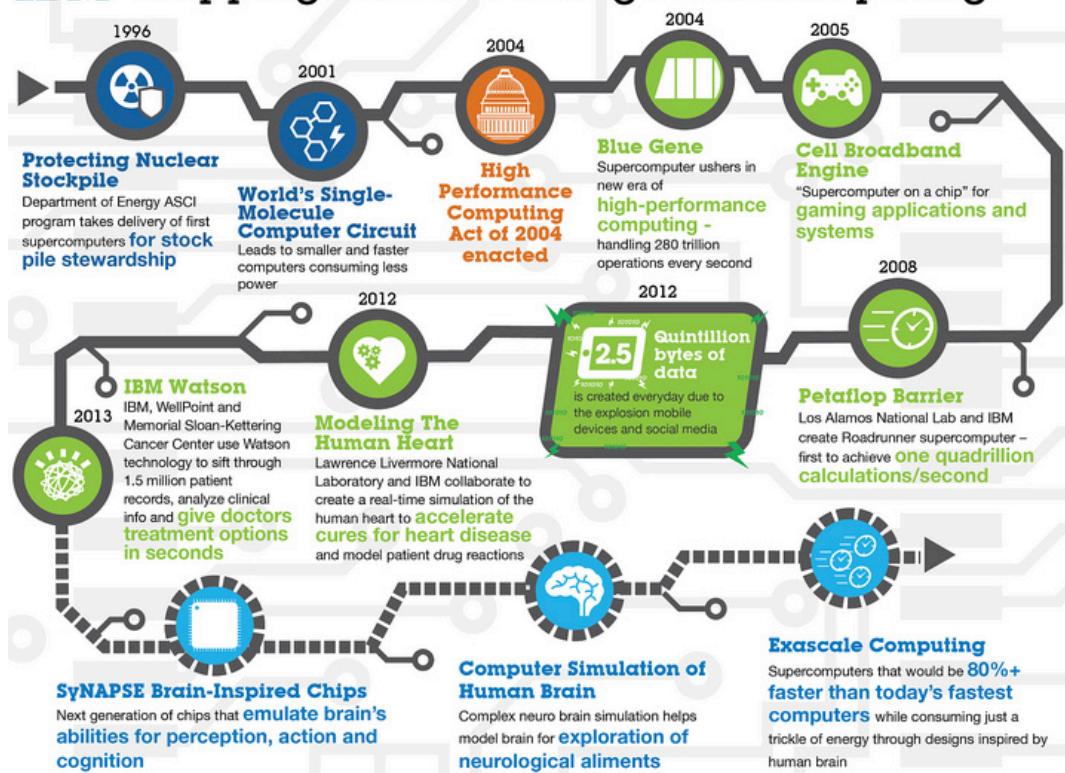
[108] Scientists can train a chimp or human to control objects purely by thought and control prosthetics that also provide feedback directly to the brain. The 'Walk Again Project' aims to produce a BCI-controlled suit enabling a quadriplegic to deliver an opening kick of the ball at the 2014 World Cup. Nicolelis has already demonstrated 'neuro-prosthesis' - a device that gives rats an extra infrared sense: 'adult rats can learn to perceive otherwise invisible infrared light through a neuroprosthesis that couples the output of a head-mounted infrared sensor to their somatosensory cortex (S1) *via* intracortical microstimulation'. DARPA's 2013 budget includes 'the Avatar program': 'Key advancements in telepresence and remote operation of ground systems are being made towards the ultimate goal of developing remotely operated robotic systems that can operate in dismounted environments… The Avatar program will develop interfaces and algorithms to enable a soldier to effectively partner with a semi-autonomous bi-pedal machine and allow it to act as the soldier's surrogate. Once developed, Avatar will allow soldiers to remain out of harm's way while still leveraging their experience and strengths to complete important missions such as sentry/perimeter control, room clearing, combat casualty recovery, and, eventually, dismounted combat maneuver' (DARPA). IARPA's ICARUS program explores 'neuroscience-based cognitive models of sensemaking … whose functional architecture conforms closely to that of the human brain.' Its TRUST program explores how to use and enhance evolved abilities to detect deception.

[109] E.g. We are understanding the causes of psychopathy: a combination of the MAOA gene, the brain desensitised to serotonin *in utero*, a major 3D pre-pubescent violent experience, and brain damage detectable with fMRI.

For example, the most successful government technology developer, DARPA, has made robotics and machine intelligence a priority with projects such as the SyNAPSE Project (led by IBM's Modha) to create 'a brain inspired electronic "chip" that mimics that function, size, and power consumption of a biological cortex.'[110] They have recently announced progress in building this new architecture, fundamentally different to the standard 'von Neumann architecture' of all normal computers.[111]

> 'We seek nothing less than to discover, demonstrate, and deliver the core algorithms of the brain and gain a deep scientific understanding of how the mind perceives, thinks, and acts. This will lead to novel cognitive systems, computing architectures, programming paradigms, practical applications, and intelligent business machines. Modern computing posits a stored program model, traditionally implemented in digital, synchronous, serial, centralized, fast, hardwired, general purpose, brittle circuits, with explicit memory-addressing imposing a dichotomy between computation and data. In stark contrast, the brain uses replicated computational units, neurons and synapses, implemented in mixed-mode analog-digital, asynchronous, parallel, distributed, slow, reconfigurable, specialized, fault-tolerant biological substrates, with implicit memory addressing blurring the boundary between computation and data. It is therefore no surprise that one cannot emulate the function, power, volume, and real-time performance of the brain within the modern computer architecture. This task requires a radically novel architecture.' Modha (IEEE,11/12).

IBM Mapping the Path to Cognitive Computing

---
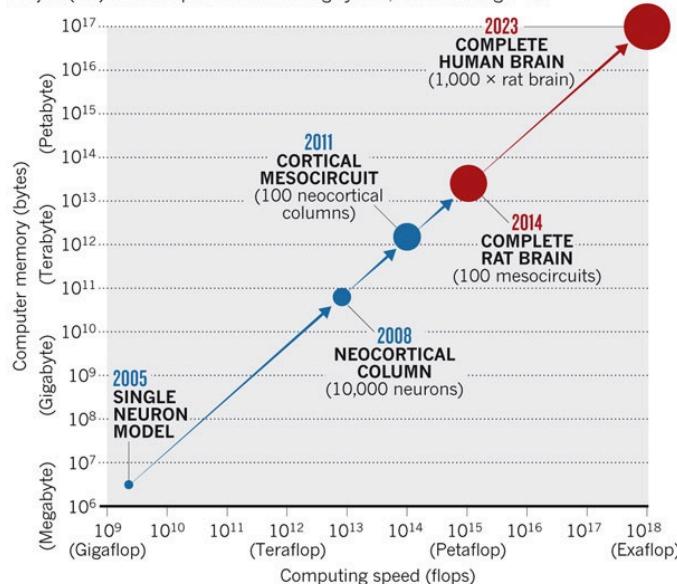
[110] 'A straightforward extrapolation of the resources needed to create a real-time human brain scale simulation shows we need about 1 to 10 exaflops [10^{18}-10^{19} flops] with 4 petabytes [4×10^{15} bytes] of memory… [T]he most optimistic current predictions for exascale computers in 2020 envision a power consumption of - at best - 20 to 30 megawatts. By contrast, the human brain takes about 20 watts. So, even under best assumptions in 2020, our brain will still be a million times more power efficient' (Horst Simon, Deputy Director of Lawrence Berkeley National Laboratory).

[111] A 5 minute Youtube; an 11/2012 update of progress. Cf. a recent update on 'memcomputing' - a fundamental shift from 'von Neumman architecture' to a neuromorphic architecture that, like neurons, processes and stores information in the same physical platform.

[The Human Brain Project](#) is aiming to model a human brain and in 2012 was awarded €1 billion by the EU. In 2005, a single neuron was simulated; in 2008, a cortical column (10^4 neurons); in 2011, 100 columns (10^6 neurons). Markram plans for a full rodent brain simulation in 2014 and a human brain simulation 2020-25.

FAR TO GO

The Blue Brain Project has steadily increased the scale of its cortical simulations through the use of cutting-edge supercomputers and ever-increasing memory resources. But the full-scale simulation called for in the proposed Human Brain Project (red) would require resources roughly 100,000 times larger still.

However, so far SyNAPSE and The Human Brain Project have not demonstrated how simulations connect to observable behaviours. In November 2012, Eliasmith *et al* (*[Science, 30/12/2012](#)*) published details of a large-scale computational model of the brain (*Semantic Pointer Architecture Unified Network, or 'Spaun'*) intended to bridge 'the brain-behaviour gap' by simulating complex behaviour. Spaun is a 'spiking neuron model' of 2.5m simulated neurons organised into subsystems resembling different brain areas. All inputs are images of characters; all outputs are movements of a physically modeled arm. Incoming visual images are compressed by successive layers of the network that extract increasingly complex information, and simple internally generated commands (e.g. 'draw a number') are 'expanded' into complex mechanical movements. Spaun simulates the working memory and an 'action selection system'. It performs eight tasks such as image recognition, copy drawing, reinforcement learning, and a fluid reasoning task 'isomorphic to the induction problems from the Raven's Progressive Matrices (RPM) test for fluid intelligence'. Spaun managed to pass some basic aspects of an IQ test. Spaun is *not* task-specific so the model could be extended to other tasks and scaled-up in other ways.

In 2013, Alex Wissner-Gross, a Harvard computer scientist, published a [paper in Physical Review](#) in 'an attempt to describe intelligence as a fundamentally thermodynamic process', proposing that intelligence can spontaneously emerge from the attempt to maximise freedom of action in the future. He built a software programme, 'ENTROPICA', designed to maximise the production of long-term entropy of any system it finds itself in. ENTROPICA then solved various problems including intelligence tests, playing games, social cooperation, trading financial instruments, and 'balancing' a physical system and so on.

'We were actually able to successfully reproduce standard intelligence tests and other cognitive behaviors, all without assigning any explicit goals…

'Think of games like chess or Go in which good players try to preserve as much freedom of action as possible. When the best computer programs play Go, they rely on a principle in which the best move is the one which preserves the greatest fraction of possible wins. When computers are equipped with this simple strategy - along with some pruning for efficiency - they begin to approach the level of Go grandmasters…

'Our causal entropy maximization theory predicts that AIs may be fundamentally antithetical to being boxed. If intelligence is a phenomenon that spontaneously emerges through causal entropy maximization, then it might mean that you could effectively reframe the entire definition of Artificial General Intelligence to be a physical effect resulting from a process that tries to avoid being boxed…

'The conventional storyline has been that we would first build a really intelligent machine, and *then* it would spontaneously decide to take over the world… We may have gotten the order of dependence all wrong. Intelligence and superintelligence may actually emerge from the effort of trying to take control of the world - and specifically, all possible futures - rather than taking control of the world being a behavior that spontaneously emerges from having superhuman machine intelligence…

'The recursive self-improving of an AI can be seen as implicitly inducing a flow over the entire space of possible AI programs. In that context, if you look at that flow over AI program space, it is conceivable that causal entropy maximization might represent a fixed point and that a recursively self-improving AI will tend to self-modify so as to do a better and better job of maximizing its future possibilities.

'In the problem solving example, I show that cooperation can emerge as a means for the systems to maximize their causal entropy, so it doesn't always have to be competition. If more future possibilities are gained through cooperation rather than competition, then cooperation by itself should spontaneously emerge, speaking to the potential for friendliness.' (Interview.)

How sophisticated will automated reasoning, perhaps merged with BCIs, become? Will we build systems that 'bootstrap' ever-better understanding in fields such as automated mathematical proof verification?

*

6. The scientific method, education, and training.

Can we create the tools needed for Nielsen's vision of decentralised coordination of expert attention and data-driven intelligence? Can we make large breakthroughs in the quality of education particularly to spread basic mathematical and scientific understanding and problem-solving skills beyond a small fraction? Can we improve the training of decision-makers?

Science

'Given enough eyeballs, all bugs are shallow,' Eric Raymond, reflecting on the phenomenal success of the 'open source' Linux operating system.

As Michael Nielsen has described in his brilliant [recent book](#), the entire scientific process itself is undergoing a dramatic revolution. In the 17th Century, the modern scientific system, in which researchers are rewarded for publishing in journals, began to evolve from the world in which Newton, Galileo *et al* kept discoveries secret and communicated in codes. This new system developed and spread over the next 300 years but is now undergoing another revolution. These changes not only promise a new era of scientific discovery - they also make integrative thinking more important and point to the sort of tools that make it easier.

Nielsen describes a future in which all scientific knowledge - data, computer code, scientific reasoning, video, libraries and more - will be on the internet, machine-readable and 'tightly integrated with a scientific social web that directs scientists' attention where it is most valuable'. First, new software tools will allow unprecedented collaboration among scientists (as with the revolutionary 'open source software' world) enabling the *decentralised coordination of expert attention* - the fundamental resource for creative problem-solving - to where it is most valuable. The restructuring of expert attention will enable a transition from the haphazard serendipity of '[Grossman moments](#)' to 'designed serendipity'.[112] Second, '*data-driven intelligence*' - a combination of large online databases plus statistical searches for patterns plus automated reasoning - allows us to find patterns in complex systems that have hitherto baffled us, though the scale of data now being collected also requires new approaches.[113]

For example, *Medline* allows searches of disease databases. *Google Flu Trends* has revealed patterns of infection. The *Sloan Digital Sky Survey* and *Ocean Observatories Initiative* enable searches of the heavens and oceans. The *Allen Brain Atlas* and similar projects will allow searches of neuronal databases. The *OpenWetware* project at MIT allows searches of genetic databases. *Google Translate* allows searches of language databases and has revolutionised translation. DARPA's '[Mind's Eye](#)' and

---

[112] E.g. The global collaboration to [play Kasparov](#) shocked chess players. [FoldIt](#) shocked biologists. [Galaxy Zoo](#) (and '[*Galaxy Zoo 2*](#)' (2013)) shocked astronomers. [Eyewire](#) may shock neuroscientists and [Phylo](#) geneticists. Tim Gowers' [Polymath project](#) shocked the maths/scientific world; Gowers said that the Polymath method was 'to normal research as driving is to pushing a car'. ('[*The Stacks Project*](#)' is an online collaborative textbook for algebraic geometry.) In 2010 a claim to have solved 'P=NP?' was analysed and disproved by [leading figures](#) in the maths/[computer science](#) communities online. In 2013, mathematicians, using the internet, rapidly improved on Zhang's weak proof of the 'twin prime conjecture' (which we have been trying to prove for >2,000 years) to lower the upper bound from 70 million to thousands in days (cf. [Polymath](#)). Could the Polymath approach be scaled up to massive networked collaborations tackling our biggest challenges, such as [the Riemann Hypothesis](#)?

[113] 'The Atlantic Ocean is roughly 350 million cubic kilometers in volume, or nearly 100 billion, billon gallons of water. If each gallon of water represented a byte or character, the Atlantic Ocean would be able to store, just barely, all the data generated by the world in 2010. Looking for a specific message or page in a document would be the equivalent of searching the Atlantic Ocean for a single 55-gallon drum barrel' (DARPA Acting Director). Twitter produces >12TB (1.2×10^{13}) and Facebook >20TB (2×10^{13}) of data per day. A six-hour Boeing flight generates ~240TB (2.4×10^{14}). Wal-Mart generates >2PB (2×10^{15}) of customer data every hour. (Source: http://arxiv.org/pdf/1301.0159v1.pdf)

XDATA programs (surveillance) may shock anyone trying to hide.[114] These statistical models are unlike the conventional scientific explanations of Newton's mechanics, Maxwell's electromagnetism or the austere equations of General Relativity or the Standard Model. They do not provide precise predictions about the behaviour of individual agents; they describe statistical relationships. They are nevertheless increasingly effective and it is possible that they are the best models we will have, for decades anyway, of extremely complex systems.

As discussed above, many problems have solution landscapes that are so vast that they cannot be exhaustively searched with any possible future technology or are theoretically infinite: for example, the number of possible 2-megabyte photos is $256^{2,000,000}$. However, for many problems once we have a billion or so examples 'we essentially have a closed set that … approximates what we need' for a meaningful statistical analysis; algorithmic analysis of photos is viable using statistical tools despite the practically infinite search landscape. Statistical speech recognition and translation have been successful because 'a large training set of the input-output behavior that we seek to automate is available to us *in the wild*' (Norvig *et al*) - that is, the translations done by humans are available on the web. These can be analysed statistically for patterns and this approach has brought great breakthroughs, without requiring the detailed specific understanding of a particular language that many in the AI community thought for decades would be required.[115]

There is widespread criticism of British science policy and funding institutions (e.g. pure maths has been badly hit by recent decisions). The institutions, priorities, and methods need a fundamental rethink and then considerably more money should be spent (saved from the closure of existing UK government functions and sensible prioritising). Britain should take a lead in developing the tools Nielsen discusses, build new institutions (such as a UK 'DARPA' focused on non-military goals, and trans-disciplinary networks), and experiment with large collaborations to solve the most important problems, such as the Millennium Problems and those discussed in this essay. Scientists could have *much* more influence on Government if they organised better.

Education

'*During the three years which I spent at Cambridge my time was wasted, as far as academical studies were concerned... I attempted mathematics, and even went during the summer of 1828 with a private tutor ... but I got on very slowly. The work was repugnant to me, chiefly from my not being able to see any meaning in the early steps in algebra. This impatience was very foolish, and in after years I have deeply regretted that I did not proceed far enough at least to understand something of the great leading principles of mathematics, for men thus endowed seem to have an extra sense.' Darwin, 1876.*

Feynman, Judging Books by Their Covers, on his experience of being asked to rate maths and science school textbooks, 1964/5: 'Everything was written by somebody who didn't know what the hell he was talking about, so it was a little bit wrong, always! And how we are going to teach well by using books written by people who don't quite understand what they're talking about, I cannot understand. I don't know why, but the books are lousy; UNIVERSALLY LOUSY!… They said things that were useless, mixed-up, ambiguous, confusing, and partially incorrect…'

'[Steve] Jobs also attacked America's education system, saying that it was hopelessly antiquated and crippled by union work rules. Until the teachers' unions were broken, there was almost no hope for education reform. Teachers should be treated as professionals, he said, not as industrial assembly-line

---

[114] DARPA's 1.8 gigapixel ARGUS drone can (as of 2013) scan 25 square miles of ground surface from 20,000 feet in enough detail to track all moving objects automatically, storing a million terabytes per day for future replays. The plan is to analyse the data from such sensors with software like Mind's Eye.

[115] Cf. '*The unreasonable effectiveness of data*'.

workers. Principals should be able to hire and fire them based on how good they were. Schools should be staying open until at least 6 p.m. and be in session eleven months of the year. It was absurd that American classrooms were still based on teachers standing at a board and using textbooks. All books, learning materials, and assessments should be digital and interactive, tailored to each student and providing feedback in real time.' What Steve Jobs told Obama shortly before his death (p. 554, Isaacson, 2012).

The education of the majority even in rich countries is between awful and mediocre. A tiny number, less than 1 percent, are educated in the basics of how the 'unreasonable effectiveness of mathematics' provides the 'language of nature' and a foundation for our scientific civilisation[116] and only a small subset of that <1% then study trans-disciplinary issues concerning the understanding, prediction and control of complex nonlinear systems. Unavoidably, the level of one's mathematical understanding imposes limits on the depth to which one can explore many subjects. For example, it is impossible to follow academic debates about IQ unless one knows roughly what 'normal distribution' and 'standard deviation' mean, and many political decisions, concerning issues such as risk, cannot be wisely taken without at least knowing of the existence of mathematical tools such as conditional probability. Only a few aspects of this problem will be mentioned.

There is widespread dishonesty about standards in English schools,[117] low aspiration even for the brightest children,[118] and a common view that only a small fraction of the population, a subset of the most able, should be given a reasonably advanced mathematical and scientific education, while many other able pupils leave school with little more than basic numeracy and some scattered, soon-forgotten facts. A reasonable overall conclusion from international comparisons, many studies, and how universities have behaved, is that overall standards have roughly stagnated over the past thirty years (at best), there are fewer awful schools, the sharp rises in GCSE results reflect easier exams rather than real educational improvements, and the skills expected of the top 20 percent of the ability range studying core A Level subjects significantly declined (while private schools continued to teach *beyond* A Levels), hence private schools have continued to dominate Oxbridge entry while even the best universities have had to change degree courses substantially.

---

[116] 'Mathematics is the queen of the sciences and the theory of numbers is the queen of mathematics.' Gauss. The EPSRC's International Review of Mathematical Sciences (2011) and the Institute of Mathematics' case studies summarise many of maths' contributions to learning and the economy.

[117] Large improvements in state-controlled test scores have not been matched by independent tests. Cf. Durham University's comparison of SAT scores at 11 and independent tests, and similar studies of GCSEs. The GCSE grades of pupils who sat PISA in 2009 were *higher* than those of pupils who got the *same* PISA maths score in 2006. A major study found a *significant decline* in the algebra and ratio skills of 14 year-olds between 1979 and 2009: <1/5 of 14 year-olds can write 11/10 as a decimal. Pupils in the top tier of countries (Shanghai, Singapore, *et al*) are about 1-2 *years* ahead of English pupils in the PISA maths. TIMSS (a more curriculum-focused international maths test) also shows England behind at primary school and the gap widening with age. A 2011 BIS survey of adults found that only ~1/5 operate in maths at a level of a GCSE 'C' or better. Research from National Numeracy (February 2012) showed that ~½ of 16-65 year olds have at best the mathematical skills of an 11 year-old. Before the EBac, only ~1/8 of those at a non-selective state school got a C or better GCSE grade in English, Maths, two sciences, history or geography, and a language (~½ of privately educated pupils do)..GCSEs are poor: cf. the Royal Society's 2011 study of Science GCSEs, Ofqual's April 2012 report, and damning analysis from the Institute of Physics and Royal Society of Chemistry. Shayer *et al* (2007) found that performance in a test of basic scientific concepts fell between 1976 and 2003. Many University courses, including Cambridge sciences, had to lengthen from the 1990s to compensate for the general devaluation of exams. Foreign languages are in crisis even in Oxbridge: forty years ago, interviews were conducted in the language, now tutors are happy if an applicant has read anything in the language. For a recent survey, cf. Coe (2013).

[118] Only 1.7% of 15-year-olds in England achieved Level 6 in PISA 2009 maths tests compared with: Shanghai (27%), Singapore (16%), Hong Kong (11%), South Korea and Switzerland (8%), Japan (6%), Finland and Germany (5%).If one looks at Levels 5 and 6, only 10% in England reach this level compared with: Shanghai (50%), Singapore (36%), Hong Kong (31%), South Korea (26%), Switzerland (24%), Finland (22%), Japan (21%). Given that those from independent schools score >50 points higher than those from maintained schools, the tiny English 1.7% may include a large over-representation of independent schools and the performance of pupils in non-grammar state schools may be worse than these figures suggest.

There is hostility to treating education as a field for objective scientific research to identify what different methods and resources might achieve for different sorts of pupils. The quality of much education research is poor. Randomised control trials (RCTs) are rarely used to evaluate programmes costing huge amounts of money. They were resisted by the medical community for decades ('don't challenge my expertise with data')[119] and this attitude still pervades education. There are many 'studies' that one cannot rely on and which have not been replicated. Methods are often based on technological constraints of centuries ago, such as lectures. Square wheels are repeatedly reinvented despite the availability of exceptional materials and subject experts are routinely ignored by professional 'educationalists'.[120] There is approximately zero connection between a) debates in Westminster and the media about education and b) relevant science, and little desire to make such connections or build the systems necessary; almost everybody prefers the current approach despite occasional talk of 'evidence-based policy' (this problem is one of the reasons we asked Ben Goldacre to review the DfE's analysis division). The political implications of discussing the effects of evolutionary influences on the variance of various characteristics (such as intelligence ('g') and conscientiousness) and the gaps between work done by natural scientists and much 'social science' commentary have also prevented rational public discussion (cf. Endnote on IQ).[121]

Westminster and Whitehall have distorted incentives to learn and improve,[122] have simultaneously taken control of curricula and exams and undermined the credibility of both, and have then blamed universities for the failures of state schools[123] and put enormous pressure on Universities and academics not to speak publicly about problems with exams, which has made rational discussion of exams impossible. Most people with power in the education system are more worried about being accused of 'elitism' (and 'dividing children into sheep and goats') than they are about problems caused by poor teaching and exams and they would rather live with those problems than deal with those accusations.[124]

---

[119] When Semmelweis, based on data, argued in the 1840s that doctors should wash their hands between autopsy rooms and childbirth, he was ridiculed, ignored, and fired.

[120] E.g. Many involved in teacher training are hostile to Cambridge University's Maths and Physics 16-18 projects (see below) and similar ventures.

[121] During my involvement in education policy 2007-12, I never come across a single person in 'the education world' who raised the work of Robert Plomin and others on IQ, genetics and schools, and whenever I raised it people would either ignore it or say something like 'well *obviously* IQ and genetics has no place in education discussions'. I therefore invited Plomin into the DfE to explain the science of IQ and genetics to officials and Ministers. Cf. this interview.

[122] E.g. The 'equivalents' racket whereby courses were deemed 'equivalent' to Maths GCSE, encouraging schools to game league tables by entering thousands for courses (created to fit the rules) of little, no, or negative value. This has largely been removed, against the predictable opposition of many powerful interests who found the racket a useful way of demonstrating 'success' that justified more funding.

[123] ~1/7 Free School Meal pupils and ~9/10 private school pupils go to University. Private school pupils (7% of all pupils) provide ~½ of places at Oxford and Cambridge mainly because of state schools' problems; this 2012 study of Oxford entrants' exam performance found that, *contra* assumptions, 'male or independent school applicants face effectively higher admission thresholds'.

[124] E.g. Almost everybody the DfE consulted 2011-13 about curriculum and exam reform was much more concerned about accusations of elitism than about the lack of ambition for the top 20%. Although they would not put it like this, most prominent people in the education world tacitly accept that failing to develop the talents of the most able is a price worth paying to be able to pose as defenders of 'equality'. The insistence that ~95% of pupils be able to take the same exam at 16 means (if one assumes symmetrical exclusions) that the exam must embrace plus and minus two standard deviations on the cognitive ability range: i.e. they exclude only the bottom 2.5% (i.e. an IQ of <~70) and top 2.5% (i.e an IQ of >~130, which is the average Physics PhD).

Chart: Comparison of English performance in international surveys versus GCSE scores (Coe)
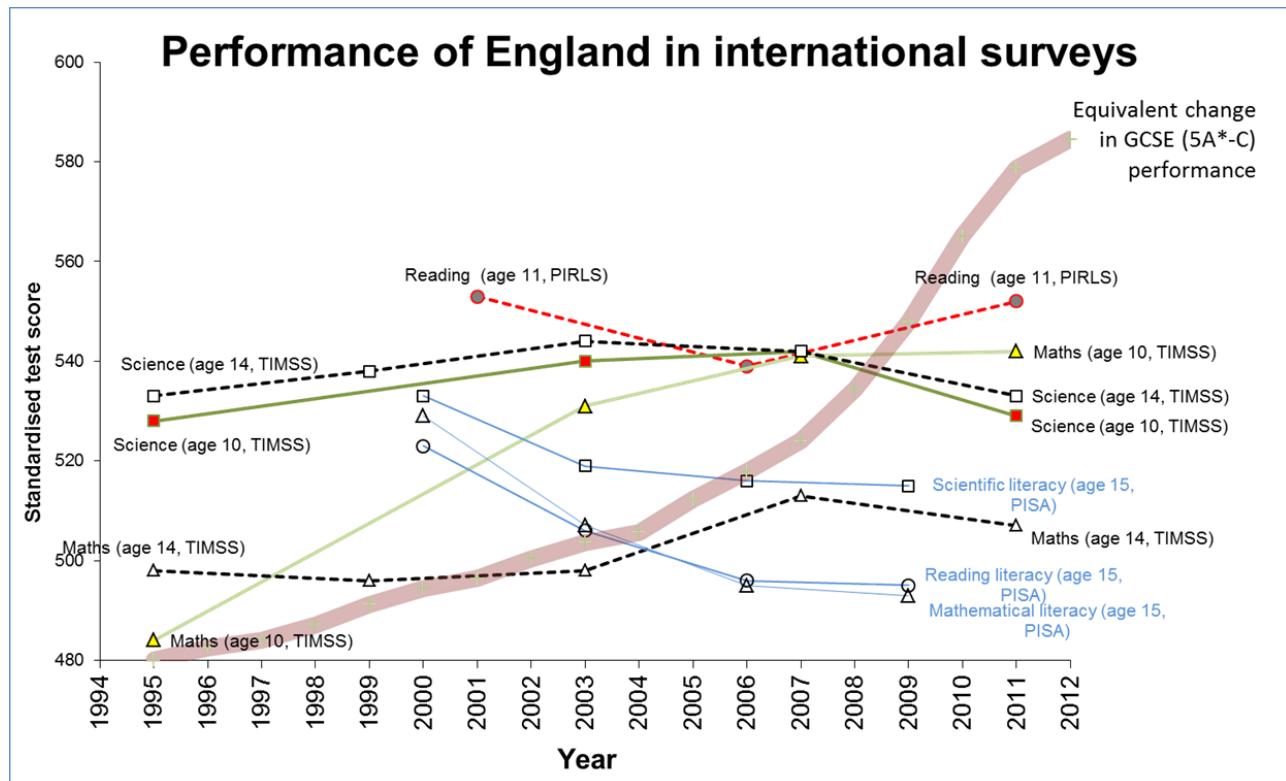
Chart: Performance on ICCAMS / CSMS Maths tests (Coe)
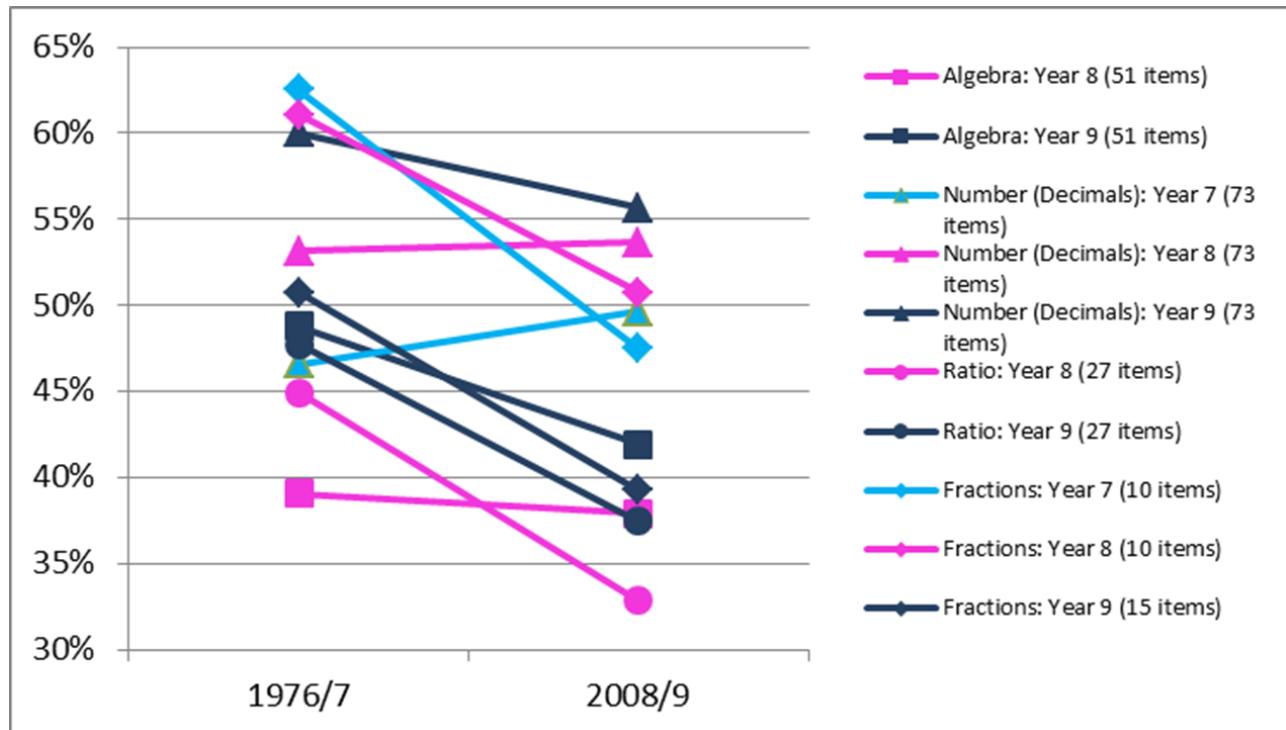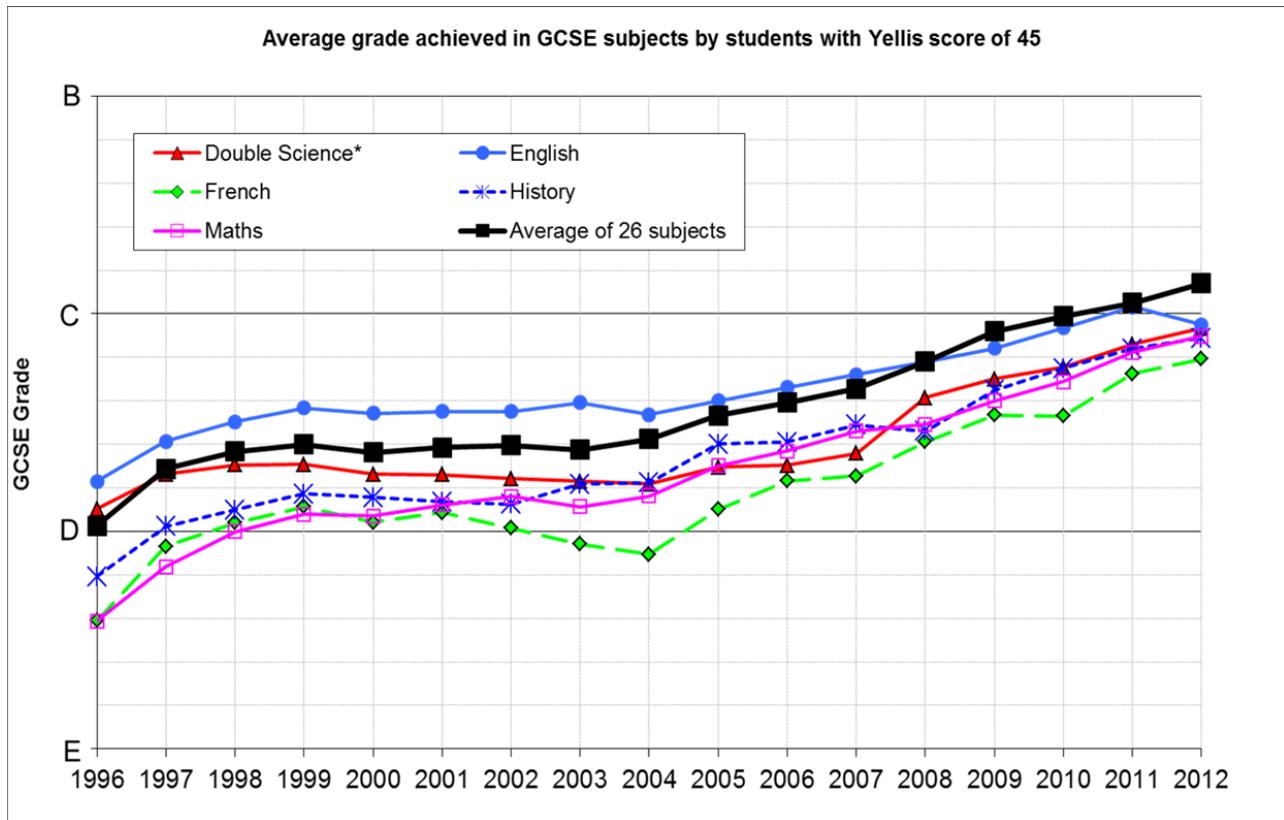
Chart: Changes in GCSE grades achieved by candidates with same maths & vocab scores each year (Coe)

Average grade achieved in GCSE subjects by students with Yellis score of 45

Legend:
- Double Science*
- French
- Maths
- English
- History
- Average of 26 subjects

(Y-axis: GCSE Grade, from E to B; X-axis: 1996 to 2012)

There is huge variation in school performance (on exams that are sub-optimal) among schools with the poorest children. About a quarter of primaries have over a quarter of their pupils leave each year who are not properly prepared for basic secondary studies and few such pupils enjoy a turnaround at secondary;[125] other primaries (including in the poorest areas) have <5% in such a desperate situation. Consider a basic benchmark: getting four-fifths of pupils to at least a 'C' in existing English and Maths GCSE. A small minority of state schools achieve this, while others with similar funding and similarly impoverished pupils struggle to get two-fifths to this level. This wide variety in performance combined with severe limits on direct parent choice means the system is partly rationed by house price.[126]

This wide variety in performance also strongly suggests that the block to achieving this basic benchmark is the management and quality of teaching in the school; the block is not poverty,[127]

---

[125] Correlation between test scores at 11 and GCSEs is strong. Of those in state schools getting *below* Level 4 in KS2 tests (~1/6 of cohort), only ~5% go on to get 5 A*-C GCSEs including English and Maths, compared to ~95% of those getting *above* Level 4 (~1/3 of cohort). If you leave state primary outside the top ¼ in maths, you have <~5% chance of getting an A in GCSE.

[126] '[T]he ability to send your children to one of the best schools rather than one of the worst will add 15 to 20 per cent to the value of your home [in England]. (Internationally, estimates range from about 5 per cent to 40 per cent.) Relative to private school this is pretty good value: it corresponds to around £28,000 or so. Since one would presumably get this money back eventually, when selling the house, the real cost is just the extra interest on a larger mortgage (or the foregone opportunity to earn interest on savings).' This would be ~£1,400 per year (assuming an interest rate of 5%); ~¼ of the amount spent by the state on each child. (Tim Harford blog)

[127] The gap between ethnic Chinese and ethnic white achievement of 5 A*-C GCSEs including English and maths is huge: in maintained schools in 2011, Chinese non-FSM pupils (79%), FSM Chinese (74%); non-FSM whites (62%), FSM whites (29%). I.e FSM Chinese greatly outperform *non-FSM* whites so poverty clearly cannot explain widespread failure (DfE, 2012). While there is some speculation that ethnic Chinese, like Jews, may have an evolved genetic IQ advantage, this could not be responsible for such a huge gap in performance.

IQ, money,[128] lack of innovation, or a lack of understanding about how to teach basics. Making a transition to a school system in which ~4/5 meet this basic level is therefore an issue of *doing things we already know how to do*; the obstacles are political and bureaucratic (such as replacing management and bad teachers despite political resistance and legal complexity), although this must not blind us to the fact that most variation in performance is due to *within* school factors (including genetics) rather than *between* school factors (see below).

There are various problems with maths and science education…

The Royal Society estimates that ~300,000 per year need some sort of post-GCSE Maths course but only ~100,000 do one now. About 6/10 now get at least a C in English and Maths GCSE; most never do any more maths after GCSE.[129] There is no widely respected 'maths for non-maths specialists' 16-18 course (see below).[130] About 70-80,000 (~1/10 of the cohort) do Maths A Level each year (of these ~⅓ come from private schools and grammars)[131] and ~1-2% also do Further Maths. In the last year for which we have data, ~0.5% (3,580 pupils) went on to get A or A* in each of A Level Maths, Further Maths, and Physics.[132] Further, many universities only demand GCSE Maths as a condition of entry even for scientific degrees, so ~20% of HE Engineering entrants, ~40% of Chemistry and Economics entrants, and ~60-70% of Biology and Computer Science entrants do *not* have A Level Maths. Less than10% of undergraduate bioscience degree courses demand A Level Maths.

Because of how courses have been devised, ~4/5 pupils leave England's schools without basic knowledge of subjects like logarithms and exponential functions which are fundamental to many theoretical and practical problems (such as compound interest and interpreting a simple chart on a log scale), and unaware of the maths and physics of Newton (basic calculus and mechanics). Less than one in ten has a grasp of the maths of probability developed in the 19th Century such as 'normal distributions' and the Central Limit Theorem ('bell curves') and conditional probability.[133] Only the 1-2% doing Further Maths study complex numbers, matrices and basic linear algebra. Basic logic and set theory (developed c. 1850-1940) do not feature in Maths or Further Maths A levels, so almost nobody leaves school with even a vague idea of the modern axiomatic approach to maths unless they go to a very unusual school or teach themselves.

---

[128] '[E]xpenditure per student has increased by 68 per cent over the past decade and yet [PISA] results have remained flat. More generally, spending per student explains less than a fifth of the performance differences between countries.' (Schleicher, 2012.)

[129] ~70% of those with an A*, ~⅓ with an A, ~5% with a B, and ~0% with a C, go on to A Level Maths.

[130] The 'Use of Maths' A Level has been strongly criticised by leading mathematicians. E.g. Two blogs (1,2) by Professor Tim Gowers (Cambridge). The current A Level Maths and Further Maths splits material into modules on a basis that, even if one accepts the need for a modular course, is suboptimal both from the overall perspective of A Level students and from the perspective of creating an AS appropriate for those *not* going on to do a maths or maths-based degree. AS does not provide a teacher with a way to select modules that cover certain basics, and is more 'half an A Level' than 'an introduction to some fundamental maths concepts that will help understanding of a wide variety of problems'.

[131] To this 1/10 can be added AS numbers, but even at 1/5 England is an international outlier for maths 16-18. Cf. '*Is the UK an Outlier?*' (Nuffield, 2010).

[132] Of that ~0.5%, ~⅓ did their GCSEs in private schools (DfE, 2012).

[133] Gigerenzer's '*Reckoning With Risk*' has terrifying stats on the inability of trained medical professionals making life and death decisions to understand the basics of conditional probability, which is not covered in the pre-16 curriculum (cf. Endnote). Current A Level modules have conditional probability and normal distributions in S1 and S2 (not compulsory), so one could have an A* in A Level Maths and Further Maths without knowing what these are. Data on who does which modules is not published by exam boards.

Crease's '*The Great Equations*' features ten equations stretching from Pythagoras through Newton's *f=ma* and gravity equation, Euler's $e^{i\pi} + 1 = 0$, Maxwell's equations for electromagnetism, Boltzmann's equation for entropy, Planck's equation for the Second Law of Thermodynamics, Einstein's equations for Special and General Relativity, Schrödinger's equation for the wave function and Heisenberg's equation for the Uncertainty Principle.[134] The vast majority of pupils by 16 have been introduced to the first two or three of Crease's equations at best, and four-fifths will have explored nothing more over the next two years including many very motivated, interested, and talented children. Only ~1% at best would have the maths to understand something of the other topics, and the education of even this 1% is increasingly criticised by top universities.[135] Various studies have concluded that we are producing only a fraction of the mathematically educated population universities and enterprises need.[136]

What to do?

The education world generally resists fiercely the idea that a large fraction of children can or should be introduced to advanced ideas *but we could substantially raise expectations without embracing 'Ender's Game'*. It is almost never asked: *how could we explore rigorously how ambitious it is realistic to be?* If you ask, including in the Royal Society, 'what proportion of kids with an IQ of X could master integration given a great teacher?', you will get only blank looks and 'I don't think anyone has researched that'. Given the lack of empirical research into what pupils with different levels of cognitive ability are capable of with good teachers, *research that obviously should be undertaken*, and given excellent schools (private or state) show high performance is possible, it is important to err on the side of over-ambition rather than continue the current low expectations. Programmes in America have shown that 'adolescents scoring 500 or higher on SAT-M or SAT-V by age 13 (top 1 in 200), can assimilate a full high school course (e.g., chemistry, English, and mathematics) in three weeks at summer residential programs for intellectually precocious youth; yet, those scoring 700 or more (top 1 in 10,000), can assimilate at least twice this amount…' (Lubinski, 2010). (See Endnote.)

---

[134] In his famous paper, Heisenberg used *Ungenauigkeit* which means 'inexactness' and also *Unbestimmtheit* meaning 'indeterminacy'. Only in the endnote thanking Bohr did he use Bohr's preferred word *Unsicherheit* which means 'uncertainty'. Heisenberg continued to use the phrase 'Indeterminacy Principle' in German. (Lindley)

[135] Leading professors privately deplore a 'collapse' of mathematical and scientific preparation for advanced studies, including the curriculum content, modular structure and style of question in A Level (including among those with A/A*). Cf. the 2012 SCORE report and Nuffield Report showed serious problems with the mathematical content of A Levels; the Institute of Physics report (2011) showed strong criticism from academics of the state of physics undergraduates' mathematical knowledge; the Centre for Bioscience criticised Biology and Chemistry A Levels and the preparation of pupils for bioscience degrees ('very many lack even the basics'); Durham University's studies showed A Level grade inflation; the House of Lords 2012 'HE in STEM' report. Few of academics' real views become public because of strong pressure not to tell the truth. Here, unusually, a Fields Medallist describes some problems with Maths A Level ('of course A Levels have got easier'). Cf. this analysis of the devaluing of physics exams, particularly problem-solving, by Professor Warner. Calculus has gone from A Level Physics, which is no longer even required by Cambridge as a condition of entry; instead one can study further A Level Maths Mechanics modules. More schools are switching to the harder pre-U and more universities are requiring extra tests beyond A Level such as 'STEP' papers (~1,000 p/a). Oxbridge is not immune; some of their degrees have lengthened to deal with these problems.

[136] E.g. *Mathematical Needs,* ACME 2011, and House of Lords (2012). Our universities produce ~5,000 maths undergraduates per year and ~6,000 physics and chemistry undergraduates per year, of whom ~¾ are UK citizens (drop-out rates are high). The UK is producing too few PhDs in mathematical sciences (~400 p/a) and the pipeline from PhD to academic position 'which was already badly malfunctioning has been seriously damaged by EPSRC decisions' (The Council for the Mathematical Sciences, 2011). 'Many UK educated PhDs … are not adequately trained to be competitive on the international job market; hence a large proportion of postdocs and junior faculty consists of researchers trained outside the UK … [and] a dwindling supply of top quality home-grown talent' (Review of Mathematical Sciences, 2010). The number of physics graduates has fallen sharply over the past 40 years. Only 6% of academics think their graduates are well prepared for a masters in Computational Biology (p.8 of report).

A combination of systematically doing what we already know and applying a scientific approach to education would allow huge progress though we do not know how far we could extend the teaching of what is today considered 'advanced' material.

1. *Largely eliminate failure with the basics in primary schools…* We should be able almost to eliminate awful primaries. Reducing the number of pupils who leave primary unable to read and do basic maths from ~¼-⅓ to less than 5 percent would enable big improvements in secondaries. We know that huge performance variance can be avoided: '[T]he most striking result from Finland is not just its high average performance but that, with only 5 per cent of student performance variation between schools, every school succeeds' (Schleicher, 2012).

Most pre-school programmes, such as Head Start (USA) and Sure Start (UK), have shown no significant sustained improvement in IQ or test scores despite billions spent over years (or they show brief improvements then 'fade out'). However, IQ is only part of what determines a child's future. As stated above regarding the unique Terman study of high IQ scientists, character, such as conscientiousness, is also vital. A small number of programmes that have been rigorously studied (particularly Perry (1962) and Abecedarian (1972)) have shown small improvements in test scores and, more importantly, long-term gains in employment, wages, health, criminality and so on (though these programmes employed much better educated people than most so there are obvious question-marks about scaleability). Studies of self-control clearly show that children with poor self-control are much more likely to be poor, have serious health problems (including dangerous addictions), and be criminals. Some programmes focused on 'executive functions', such as self-control, *have succeeded* and improved measures such as employment, crime, and drugs, though it has proved hard to scale up those few successes.

Overall, there is great political pressure to spend money on such things as Sure Start but little scientific testing, refinement, and changing budgets to reinforce demonstrated success, therefore billions have been spent with no real gains. The billions now spent in Britain should be tied to 'randomised control trials' and there needs to be more experimentation with how to expand experiments that seem to work (lots of small things work because of specific individuals and therefore do not work when expanded). *If* - and it remains a big if - experiments such as Perry could be expanded, then they would pay for themselves (in lower prison, health, unemployment bills etc) and provide a much better foundation for schools but this requires major policy changes and the existing lobby groups would likely fight an empirical approach partly because it would be seen as a threat ('an insult') to existing staff.[137]

2. *Largely eliminate failure with the basics in secondary schools…* Achieving (1) would, along with other things we already know how to do, help reduce the failure rate in secondary schools such that ~4/5 reach roughly the level now defined as at least a 'C' in English and Maths 'international O Level' exams set by Cambridge Assessment - which Singapore and a minority of English state schools suggest is a reasonable goal. Similarly, learning from what we already know works in places like Switzerland,and Singapore would provide high quality vocational courses for those who do not

---

[137] Cf. *Science (2011) special issue* on early years and Endnote on IQ. '*A gradient of childhood self-control predicts health, wealth, and public safety*' (Moffitt et al, PNAS, 2/2011). '*Preschools reduce early academic-achievement gaps: a longitudinal twin approach*' (Tucker-Drob, Psychological Science, 2/2012): a unique twin-study of pre-school effects. '*Why children succeed*' by Paul Tough (2012) summarises much of the literature including James Heckman's work. Is anybody doing a twin/adoption study involving children going to high quality pre-school *then* high quality schools, such as KIPP, to see whether the widely reported 'fade-out' of pre-school gains could be maintained by excellent schools?

continue with academic studies post-14/16.[138] (This is not to say that existing or reformed GCSEs are the right way to measure performance, see below.)

Both (1) and (2) concern *stopping people getting basics wrong* and this requires a) the removal of an unknown proportion of senior management, middle management, and teachers and b) retraining of others. See below for ideas on *how* to do this.

3. *A scientific approach to teaching practice…*

Feynman's famous Caltech commencement address, '*Cargo Cult Science*' (1974) discusses the way in which much 'education research' copies only some external forms of natural sciences but misses the fundamentals.

> 'There are big schools of reading methods and mathematical methods, and so forth, but if you notice, you'll see the reading scores keep going down … *There's* a witch-doctor remedy that doesn't work… Yet these things are said to be scientific. We study them. And I think ordinary people with commonsense ideas are intimidated by this pseudoscience…
>
> 'I think the educational and psychological studies I mentioned are examples of what I would like to call Cargo Cult Science. In the South Seas there is a Cargo Cult of people. During the war they saw airplanes land with lots of good materials, and they want the same thing to happen now. So they've arranged to make things like runways, to put fires along the sides of the runways, to make a wooden hut for a man to sit in, with two wooden pieces on his head like headphones and bars of bamboo sticking out like antennas - he's the controller - and they wait for the airplanes to land. They're doing everything right. The form is perfect. It looks exactly the way it looked before. But it doesn't work. No airplanes land. So I call these things Cargo Cult Science, because they follow all the apparent precepts and forms of scientific investigation, but they're missing something essential, because the planes don't land.'[139]

In medicine, biological knowledge is translated into medical practice and in other fields there is translation from basic knowledge (e.g. physics) into practice (e.g. engineering). It is necessary to do the same in education and there are some promising signs (cf. this paper summarising some recent findings).

Because so many of those responsible for devaluing standards use some of the language of cutting edge science (e.g. 'peer collaboration'), it has provoked many of those who wish to reverse such devaluation to reject arguments that are based on science *because they sound like pseudoscience to people who have to listen to debates about schools*. E.g. 1. 'Peer collaboration' has been used to justify useless activities, *but it is also a proven technique when used by great university teachers* (see below and Endnote). 2. Some argue that '*knowledge isn't important because of the internet - critical thinking skills are more important.*' Cognitive science tells us that you need to know stuff in order to learn more stuff - knowledge in a brain is a network, and you need to master certain facts and concepts in

---

[138] Implementing Alison Wolf's 2011 Report has already improved things: e.g. by removing large numbers of poor 'equivalent' courses that wasted pupils time and sometimes had *negative* labour market value, changing the 16-19 funding system to remove perverse incentives to do poor courses, requiring the 4/10 who do not get a C in GCSE Maths to continue maths studies, improving the quality of apprenticeships (more work needed).

[139] This speech also has a good lesson: 'But there is one feature I notice that is generally missing in cargo cult science… It's a kind of scientific integrity, a principle of scientific thought that corresponds to a kind of utter honesty… Details that could throw doubt on your interpretation must be given, if you know them. You must do the best you can - if you know anything at all wrong, or possibly wrong - to explain it. If you make a theory, for example, and advertise it, or put it out, then you must also put down all the facts that disagree with it, as well as those that agree with it… The easiest way to explain this idea is to contrast it, for example, with advertising…'

order to be able to make sense of others. You cannot master *differentiation and integration* in mathematics if your brain does not already understand *algebra*, which in turn requires *arithmetic*. The more you know, the easier it is to make sense of new things. Experts do better because they 'chunk' together lots of individual things in higher level concepts - networks of abstractions - which have a lot of compressed information and allow them to make sense of new information (experts can also use their networks to piece together things they have forgotten). 'Critical thinking skills' is not an abstract discipline that can be taught in the absence of traditional subject matter such as mathematical problem-solving, understanding the scientific method, and extended writing. However, the fact that people are wrong to argue 'knowledge isn't important' *does not mean that anything with the word skill in should be ignored* (e.g. university physicists often use terms like 'crucial skills' to mean 'can use calculus to solve hard problems', which is very different than what Mick Waters means by 'skills'). Is calculus 'knowledge' or 'skill'? It is *both* and this terminology is mainly a cul-de-sac. 3. Some argue that '*rote learning is a waste of time / damaging.*' Some things, like basic words, do not have inherent meaning - the sound of 'mother' is arbitrary and just has to be learned. Further, lots of things require practice: who thinks that teaching the commutative law of addition would produce successful children *without* practice? Moving things from the *short-term* memory (which is extremely limited even for the cleverest) to the *long-term* memory needs building proteins in the brain and this needs practice (except in extremely unusual people who instantly remember things permanently). However, sometimes in reaction to 'rote learning is bad' people respond as if anything that is not a teacher lecturing the class should be banished. (Cf. Hirsch, 2006, 2013 and Endnote on Wieman and Mazur).

This 1984 study ('*The Two Sigma Problem*'[140]) randomly assigned school-age students with comparable test scores (and other characteristics) to three groups: A) Conventional Class, in which ~30 students learned a subject in a class; B) Mastery Learning, in which students learn as in A but with formative tests of understanding, feedback, and correction in the pursuit of mastery; C) Tutoring, in which 1-3 students learned the same material with a good tutor and periodic formative tests. It found that one-to-one tutoring produced learning gains for almost everyone *two standard deviations above the Conventional Class: 'the average tutored student was above 98% of the students in the control class.*' Further, for those tutored, there was a *lower correlation* of students' prior aptitude test scores and their final achievement scores, showing students of all abilities could make large improvements. This finding has been replicated.

A series of studies has explored why tutoring works (cf. '*Learning from human tutoring*', Chi *et al*, 2001). One of the main findings is that tutoring works when tutors do not just tell pupils the answers but help them, *via* lots of probing and questioning, to think through problems themselves. Students who are prompted to explain to themselves what something means score *significantly higher* than those who just read things (even multiple times). 'Learning research has shown clearly that the process of wrestling with a problem in a small-group discussion can enhance both student understanding and problem solving' (op. cit).

Further, *testing should not be thought of as separate from learning* but as integral to it. Regular testing not only gives teachers feedback but significantly improves learning more than many other widely used methods particularly if pupils routinely return to material over time to embed it in the long-term memory ('The positive effects of testing were dramatic… Testing has a powerful effect on long-term retention', Roediger, cf. Endnote). It is possible that many of the lessons from learning in tutorials could be applied to larger settings (cf. '*The Role of the Lecturer as Tutor*', 2012).

---

[140] Here is a slightly different version. It is odd that during six years involvement in education debates, I have never heard one person mention it in a meeting.

We know that there are profound problems with university science teaching *via* traditional lectures and we are making some discoveries about how to improve it. For example, Nobel-winning physicist Carl Wieman decided to test how successfully he was teaching science. To his shock he discovered that *the vast majority of his students did not understand what he was talking about.* His graduate students would arrive in his lab but after 17 years of success 'when they were given research projects to work on, *they were clueless* about how to proceed' and 'often it seemed they didn't even really understand what physics was.' However, after a few years in his lab they would be transformed into expert physicists. This was 'a consistent pattern' over many years.

First, a large number of studies shows that *students (including elite students at the best universities) remember little from traditional lectures.* This is unsurprising given 'one of the most well established - yet widely ignored - results of cognitive science: the extremely limited capacity of the short-term working memory.' Second, analysis of results shows that students following a a traditional lecture course master 'no more than about 30 percent of the key concepts that they didn't already know at the start of the course' and this finding is 'largely independent of lecturer quality, class size, and institution.' Third, many students become *more* confused after a traditional course. Fourth, we also know that a scientific approach to education and the use of technology is demonstrating advances, sometimes dramatic, though this has little influence on school debates.[141]

> 'We measured the learning of a specific set of topics and objectives when taught by 3 hours of traditional lecture given by an experienced highly rated instructor and 3 hours of instruction given by a trained but inexperienced instructor using instruction based on research in cognitive psychology and physics education. The comparison was made between two large sections (N = 267 and N = 271) of an introductory undergraduate physics course. We found increased student attendance, higher engagement, and more than twice the learning in the section taught using research-based instruction… The two distributions have little overlap, demonstrating that the differences in learning between the two sections exist for essentially the entire student population. The standard deviation calculated for both sections was about 13%, giving **an effect size for the difference between the two sections of 2.5 standard deviations**. [O]ther science and engineering classroom studies report effect sizes less than 1.0. An effect size of 2, obtained with trained personal tutors, is claimed to be the largest observed for any educational intervention… This result is likely to generalize to a variety of postsecondary courses.' ('*Improved Learning in a Large-Enrollment Physics Class*', Wieman *et al*, *Science* (2011), emphasis added.)

Another field of research that has big potential to inform education policy and improve education is *genetics* though it remains ignored in current debates outside a tiny group.[142] When it is not ignored, it is often misunderstood, both by those such as Malcolm Gladwell (who wrongly downplayed the importance of genes, cf. a review by Pinker) and by those who wish to use genetics to justify the view that 'they are doomed by their genes'. However, it is likely that public debate will have to stop ignoring the subject within the next decade (and Epstein's *The Sport Gene* may bring the issue to more prominence shortly).

---

[141] Cf. This article by Wieman and this by Mazur. A 2011 *Science* special issue summarised knowledge about early years education including the use of technology. This NYT article summarises some recent findings. (e.g testing is usually portrayed in education debates as destructive but studies suggest that it is a more powerful tool for learning than many things recommended by educationalists. 'Retrieval is not merely a read out of the knowledge stored in one's mind – the act of reconstructing knowledge itself enhances learning. This dynamic perspective on the human mind can pave the way for the design of new educational activities based on consideration of retrieval processes.' Much of Wieman's and Mazur's message is: 'A lot of educational and cognitive research can be reduced to this basic principle: People learn by creating their own understanding' (Wieman). Cf. Endnote.

[142] It is a sign of the fundamental problems with 'education research' that the Institute of Education is very hostile to research on genetics and education.

While Francis Galton concluded that talent in various fields was primarily genetic, Watson (the founder of behaviourism) claimed (1930) that he could 'guarantee to take any one ['healthy infant'] at random and train him to become any type of specialist … even beggar-man and thief, regardless of his talents.' This view has dominated ever since despite modern developments in genetics.

In 1993, Ericsson *et al* published research purporting to show that '10,000 hours of deliberate practice' is what distinguishes 'experts' in areas such as music and chess, and that 'innate ability' is not relevant.

> 'Our theoretical framework can also provide a sufficient account of the major facts about the nature and scarcity of exceptional performance. Our account does not depend on any scarcity of innate ability (talent)… [I]ndividual differences in ultimate performance can largely be accounted for by differential amounts of past and current levels of practice.'[143]

This view had a large influence on the media and various books (e.g. Gladwell's '*Outliers*', Shenk's '*The Genius in All of Us*', Syed's '*Bounce*', and Brooks' '*The Social Animal*') have promoted the idea that people require the same amount of practice regardless of talent and that '10,000 hours' is a sort of magic number - put in those hours and you too can be great, don't worry about your genes.

Recent analysis of these studies has confirmed that such claims are greatly exaggerated: 'deliberate practice' does not account for most of the variance in performance (only about a third of the variance in chess and music), and 'some people require much less deliberate practice than other people to reach an elite level of performance'.[144]

Work by one of the pioneers of behavioural genetics, Robert Plomin, has shown that most of the variation in performance of children in English schools is accounted for by *within school* factors (not *between* school factors), of which the largest factor is genes. Scores in the phonics test show ~70% heritability; scores in National Curriculum reading and maths tests at 7, 9, and 12 show ~60-70% heritability; and scores in English, Maths and Science GCSEs show ~60% heritability in a just-completed twin study (the GCSE data will be published later in 2013).[145] In contrast, the overall effects of shared environment (including all family and school influences shared by the twins growing up in the same family and attending the same school) accounts for only about a third of the variance of GCSE scores.[146] Educational achievement in school is *more heritable* than IQ in English school children: i.e the heritability of what is directly taught is *higher* than what is not directly taught. Perhaps differential performance in educational achievement is heritable *because it is taught*: that is, roughly similar schools teaching the same material reduces a major source of environmental variation, therefore the variation that remains is even more due to genetic variation. Similarly, this paper (Science, 23/4/2010) shows how good teachers improve reading standards for all but this means that *the variance that remains is more due to genetic differences*. This leads to a conclusion almost completely at odds with prevailing conventional wisdom in political

---

[143] In a later article (2007), Ericsson wrote that 'the only innate differences that turn out to be significant - and they matter primarily in sports - are height and body size.'

[144] The conscientiousness necessary to engage in thousands of hours of practice is itself a heritable characteristic. Interestingly, personality factors are not one of the 'other' things that accounts for variance in performance: 'effects of the personality factors on performance in these studies were … fully mediated through and thus completely explained by individual differences in deliberate practice. So, personality factors may explain why some people engage in more deliberate practice than others, but they do not appear to independently explain individual differences in performance.'

[145] Cf. '*Genetic influence on educational achievement*', Plomin *et al* (forthcoming). There is an interesting difference between sciences (58% heritable) and humanities (42% heritable).

[146] It is interesting how high the *shared* environmental effects are in this study given how low they are for IQ.

and academic debates over education: *differences in educational achievement are not mainly because of 'richer parents buying greater opportunity' and the successful pursuit of educational opportunity and 'social mobility' will increase heritability of educational achievement.*

According to Plomin, large-scale genome wide association studies (GWAS) are likely to discover some of the genes responsible for general cognitive ability ('g') and specific abilities and disabilities (see Section 4 above). This would create a whole new field of study and enable *truly personalised education* including early intervention for specific learning difficulties. When forced to confront such scientific developments, the education world and politicians are likely to handle them badly partly because there is such strong resistance across the political spectrum to accepting scientific evidence on genetics and partly because those who dominate debate have great, *but greatly misplaced*, confidence in their abilities to understand the statistical methods needed to discuss genetics - methods often requiring skills beyond those of most social scientists, never mind politicians. Most of those who now dominate discussions of issues such as 'social mobility' entirely ignore genetics and therefore their arguments are at best misleading and often worthless.

For example, a leading sociologist, John Goldthorpe (Oxford), recently wrote '*Understanding - and misunderstanding - social mobility in Britain*' (2012) which explained various ways in which economists have misled the media and politicians into making claims about the 'stalling' or 'decline' of social mobility in Britain. In a paper about class and wealth across generations, he ignores genetics entirely. *However, using parent-offspring correlations as an index of 'social mobility' is fundamentally flawed because the correlations are significantly genetic - not environmental.* Raising school performance of poorer children is an inherently worthwhile thing to try to do but it would not necessarily lower parent-offspring correlations (nor change heritability estimates).[147] When people look at the gaps between rich and poor children that already exist at a young age (3-5), they almost universally assume that these differences are because of environmental reasons ('privileges of wealth') and ignore genetics.

It is reasonable to hope that the combination of 1) finding the genes responsible for cognitive abilities, 2) scientific research on teaching methods, and 3) the power of computers to personalise learning will bring dramatic improvements to education - but this will not remove genetic influence over the variation in outcomes or 'close the gap between rich and poor'. 'The good school ... does not diminish individual differences; it increases them. It raises the mean and increases the variance' (Elliot Eisner, Stanford). Good schools, in the sense of 'teaching children of different natural abilities as well as possible', will not 'eliminate gaps' - they will actually increase gaps between those of different abilities, but they will also raise floors and averages and give all children the opportunity to make the most of their genetic inheritance (personality as well as IQ). Far from being a reason to be pessimistic, or to think that 'schools and education don't matter, nature will out', the scientific exploration of intelligence and learning is not only a good in itself but will help us design education policy more wisely (it may motivate people to spend *more* on the education of the less fortunate). *One can both listen to basic science on genetics and regard as a priority the improvement of state schools; accepting we are evolved creatures does not mean 'giving up on the less fortunate' (the fear of some on the Left) or 'giving up on personal responsibility' (the fear of some on the Right).*

An Endnote discusses some of the evidence about IQ, genetics, SATs and so on: for example, the Terman and Roe studies and the groundbreaking 'Study of Mathematically Precocious Youth' (SMPY), which shows that even differences within the top 1% of maths ability at age 13 predict significantly different chances of getting a doctorate, filing a patent, and so on.

---

[147] Goldthorpe points out that class shows a stronger parent-offspring correlation than income or earnings. He does not mention that class is more heritable than income.

4. *Maths for most 16-18…* It should become the norm, as it is in other advanced countries, that the majority of pupils will continue with some sort of Maths 16-18 course. If (1) and (2) are achieved, then more will be able to do A Levels but there is also an urgent need for a mathematical problem-solving course based on the ideas of Fields Medallist Tim Gowers, who has sketched [on his blog](#) how such a course could be constructed, and others ([MEI is now developing](#) such a course).[148] Such a course, similar to '*Physics for Future Presidents*' at Berkeley, *which should also be widely adopted in schools*, could have a lot of common questions and problems that are studied at different levels of mathematical sophistication (ranging from those who struggle to do basic formal maths to Further Maths A Level or STEP students).[149]

This would have various positive effects on Universities. It would also require further shifting of resources and methods in teacher training, including projects for academics with deep subject knowledge to redefine fundamentals in the subject for those who are responsible for large scale teacher training, since the current system gives almost total power over training to people who do not know, and are often hostile to, how true experts understand the field. A significant rise in the numbers studying maths 16-18 requires that Universities are much more forceful in requiring maths qualifications as a condition of entry to courses. Now, Universities complain constantly about the lack of entrants' mathematical skills and say it is not up to them to break the chicken-and-egg cycle but until Universities make both A Level and 'maths for future presidents' type courses formal entry requirements, it is unlikely there will be a major change soon.

5. *Specialist schools from which all schools (and Universities) can learn…* We know that at the top end of the ability range, specialist schools, such as the famous Russian 'Kolmogorov schools',[150] the French *prépas* such as *Louis-Le-Grand* (which half France's Fields Medallists attended), or Singapore's

---

[148] The Senior Physics Challenge, run by Cambridge University for 6th formers to test problem-solving skills, uses questions very similar to Gowers' suggestions.

[149] What sort of things should people learn? For example… *Combinations and permutations*: e.g. how many five card hands are there in a deck of 52 cards, or how many possible choices are there in a given situation? *Conditional probability*: e.g. what is the probability that you have a disease given you have tested positive and the test is x% accurate? (Gigerenzer has shown how teaching people to convert problems requiring Bayes' Theory into 'natural frequencies' greatly improves understanding of many real world problems. Cf. Endnote.) *Statistics*: e.g. what does it mean to say 'IQ is correlated with X', or 'a randomised trial shows that X works'? What are the Law of Large Numbers, 'normal distribution', and the Central Limit Theorem? *Estimation*: this vital skill is tested with pencil and paper tests in Cambridge's Senior Physics Challenge (e.g how many photons per second hit the eye from Alpha Centauri?). Different pupils should learn concepts at different levels: eg. many pupils would benefit from [this](#) on randomness and standard deviation.

[150] The network of Russian maths and physics schools initiated by Kolmogorov (one of the 20th Century's leading mathematicians) educated many faculty members of western universities (and Grigori Perelman, who solved the Poincaré Conjecture in 2003). The experience of these schools is that 1) they must take great care in selecting pupils (e.g a summer school system); 2) they need complete freedom over the curriculum and testing (Kolmogorov had students learning set theory, non-Euclidean geometry, topology, and other subjects not in a usual school curriculum; the Novosibirsk school used Feynman's famous Lectures as their textbooks); 3) ideally they are integrated in the life of a university with academics teaching maths and sciences. These schools gave me the idea to copy them in England.

equivalent,[151] as well as a handful of England's best schools and other examples,[152] show that it is possible to educate the most able and interested pupils to an extremely high level. We know (thanks to studies such as SMPY and Terman) that although cognitive tests are noisy, we can quite accurately identify the top ~2% in IQ[153] on whom we rely for scientific breakthroughs and even crude tests at 13 predict differences *within the top 1%* in future scientific achievements (*contra* repeated assertions, tests of cognitive ability are much better indicators of future academic performance than social class; cf. Endnote). We should give this ~2%[154] a specialist education as per Eton or Kolmogorov, including deep problem-solving skills in maths and physics. 'We want to educate leaders, the innovators of society. Let's turn our students into *real* problem solvers… We have to train people to tackle situations they have not encountered before. Most instructors avoid this like the plague, because the students dislike it. Even at Harvard, we tend to keep students in their comfort zone. The first step in developing those skills is stepping into unknown territory' (Mazur, Professor of Physics, Harvard). What is it like to have an understanding of advanced mathematical problem-solving skills? This eloquent essay gives a clue.

The first English specialist maths schools, run by King's College and Exeter University, have been approved by the Department for Education and will open in 2014. All of the pupils will be prepared for the maths 'STEP' paper that Cambridge requires for entry (or Oxford's equivalent) - an exam that sets challenging problems involving *unfamiliar* ways of considering familiar material, rather than the formulaic multi-step questions of A Level.[155]

It is also vital that such students be given a broad education and a feel for how to apply fundamental skills in maths and physics to many problems. They should be reading books like Gell

---

[151] The National University of Singapore High School of Mathematics and Science (a school affiliated to the NUS) offers courses that include: advanced maths (including calculus at an undergraduate level, linear algebra and matrices, group theory, number theory and set theory), natural sciences, as well as introductions to software engineering, Operational Research, game theory. Students also study design and engineering, they have a broad curriculum including arts and humanities, and they are taught (as part of the 'Da Vinci programme') general ideas about problem-solving and integrating different subjects to prepare them for a world in which new research areas rapidly evolve.

[152] E.g. The famous central European gymnasia such as the Lutheran Academy that educated the likes of von Neumann (not only did such schools do a brilliant job themselves but they also did a much better job than many modern educators in introducing exceptional children to University studies without ruining them). This page has links to essays about different countries' specialist maths schools, e.g. this fascinating anonymous essay by an old boy of *Louis-Le-Grand*. A billionaire who attended one of these famous schools told this author: 'The intellectual level at my [maths] school was here [hand raised high]; at [prestigious US university], it was here [hand lowered]; at Harvard Business School, here [hand lowered further] - people there [HBS] were good at bullshit and making money but not adding real value or solving hard problems.'

[153] Because g≈IQ is normally distributed with an average value of 100 and a standard deviation (SD) of 15, ~⅔ of the population lie between 85-115 and ~95% lie between 70-130. The top 2% in IQ are those >2SD in IQ, so >130, which is roughly the IQ of an average physics PhD (according to one study the average Sociology PhD is ~115, which is the same as the average undergraduate, and the average Education PhD is ~110). +3SD is >145 and is ~1:1,000; +4SD is >160 and is 1:30,000 (global population ~250,000); +5SD is 175 and is 1:3½ million (global population ~2,000); >6SD is >180 and is 1 in a billion (global population can fit in a SUV). A Fermi estimate suggests that ~6-700 +3SD and ~20 +4SD children start school in England each year: identifying them and giving them what they need should be a priority.

[154] A decentralised system will develop different paths and allow children to move between them - a single noisy selection test at 11 is too crude a method to decide a child should be permanently put in a 'not academic' stream. The fact that Finland is among the top handful of jurisdictions is evidence that general excellence as currently defined does not *require* selection, though exceptional performance of the Kolmogorov variety obviously does. Higher IQ pupils selected for specialist schools will get a better education than in comprehensives but there are also complicated aspects of peer-group effects to consider. A strong argument for some selection is the principle that a high/low IQ child should not have to suffer being in a comprehensive rather than in a more specialist environment because of 'peer effects' arguments, and that it is not fair that richer pupils have this option while poorer pupils do not.

[155] Other maths departments were enthusiastic about the idea but Vice Chancellor offices were hostile because of the political fear of accusations of 'elitism'. Hopefully the recent support of Les Ebdon for the idea will change this.

Mann's *The Quark and the Jaguar*, E.O. Wilson's *Consilience*, Pinker's *The Blank Slate*, Beinhocker's *Complexity Economics*, Mitchell's *Complexity*, and others in the Endnote; studying the issues in Muller's *Physics for Future Presidents*; learning how to programme agent-based models and simulations; reading classics from Thucydides to Dostoyevsky; and experiencing very different environments and responsibilities (cf. the footnote above on Singapore's NUS specialist school syllabus, and Kolmogorov stressed the importance of exposing such pupils to Shakespeare, the arts, music and so on).

These ideas should be extended to Universities.[156] Caltech requires *all* undergraduates, including humanities students, to do introductory courses in advanced mathematical and scientific concepts and this approach has proved extremely successful.[157] Over the past decade, Stanford has made a major effort to develop interdisciplinary courses and what it calls 'T-shaped students' who have deep expertise in a specialism *and* broad understanding of many fields.[158] Universities should be thinking hard about new inter-disciplinary courses: *instead of Politics, Philosophy, and Economics (which has not done a good job in preparing leaders), we need Ancient and Modern History, Physics for Future Presidents, and Applied Maths*; or, *The History of Ideas, Maths for Future Presidents, and Programming*.

How far might lessons from such schools be extended throughout the system?[159]

6. *Web-based curricula, MOOCs, and HE/FE…* We should use the internet for the decentralised construction of new curricula incorporating the world's best content. This enables much more ambition for, and quality of, education beyond tiny elites. For example, this essay by Michael Nielsen (author of the standard textbook on quantum computers) is about one of the most difficult subjects - the foundations of quantum mechanics - but he makes very difficult concepts understandable to teenagers. A Professor of Theoretical Physics at Cambridge, Mark Warner, has recently published '*A Cavendish Quantum Mechanics Primer*' to help *school children* (doing A Level Physics and Maths) understand quantum mechanics (understand in the sense of *calculate* and do the subject properly). Fields Medallist Terry Tao has posted notes on teaching Special Relativity to LA High School pupils. The Russians developed advanced school level maths textbooks. How many other examples are there that are not being used?[160]

---

[156] Various people have written about the difference in the education of the older generation of von Neumann, Einstein, Gödel *et al* (who would stroll around Princeton discussing Parmenides and Plato) and the post-war generation. 'The withdrawal of philosophy into a "professional" shell of its own has had disastrous consequences. The younger generation of physicists, the Feynmans, the Schwingers, etc., may be very bright… But they are uncivilized savages, they lack in philosophical depth' (Feyerbrand, 1969). 'Perhaps the one thing I regret [about my education] is not taking enough humanities courses in high school and college, though I think at that age I would have been too immature to benefit from them anyway; they were certainly the courses I had the most difficulty with at the time. Unlike science and mathematics, it seems the humanities require some experience with life itself in order to truly appreciate the subject' (Terry Tao, Fields Medallist, identified age 2 as exceptional, *Insights from SMPY's Greatest Child Prodigies*, 2006).

[157] Vernon Smith on his experience of Caltech and Harvard: 'The first thing to which one has to adapt is the fact that no matter how high people might sample in the right tail of the distribution for "intelligence," … that sample is still normally distributed in performing on the materials in the Caltech curriculum. The second thing you learn, if you were reared with my naive background, is the incredible arrogance that develops in conjunction with the acquisition of what you ultimately come to realize is a really very, very small bit of knowledge compared with our vast human ignorance.'

[158] E.g. Stanford's Institute of Design gives students interdisciplinary projects in health, education, and the developing world. A 2012 internal report on Stanford was, given that it is generally regarded as one of the world's top ten universities, interestingly critical of 'how little attention most departments and programs' have given to cultivating the capacity 'to forge fresh connections' between areas of knowledge and different skills, and the 'curiously compartmentalized lives' of their students.

[159] Cf. Fields Medallist Terry Tao blog about maths education.

[160] Nielsen also provides an introduction to sophisticated ideas about causation; the classic book on this is *Causality: Models, Reasoning, and Inference*.

Universities have already posted lectures online but now they are developing free *interactive* courses - massively open online courses (MOOCs) - a phenomenon described by John Hennessy, President of Stanford University, as a 'tsunami'.[161] It should also be possible to use online tests (independent of state control) to get better information about what is learned, such as the Force Concepts Inventory Test used in physics. Recently, Cambridge University Mathematics department and Physics department have, because of serious concern about the devaluation of exams (e.g thoughts on [maths A Level by Professor Gowers here](#) and an [analysis of physics exams by Professor Warner here](#)), launched projects to develop new 16-18 curriculum materials, problem sets, tests, and online learning - two projects that may have a large impact and increase the quality and numbers leaving schools with advanced skills.

HE and FE reform… We should change national funding priorities a) to extend the proven Oxbridge tutorial system beyond those two universities and b) to increase substantially funding for the top tier of universities.[162] The Government spends over half a trillion pounds per year of which vast amounts are wasted and it could easily divert a few billion pounds, *if* it could prioritise, to strengthen world class humanities, maths, and science departments (including fixing the broken PhD pipeline) and develop new strengths. Increased funding should be tied to a shift back to independence for universities; political control has corrupted admissions policy and debates on standards. It should also, as Carl Wieman has argued (*[Science, 19/4/2013](#)*), be tied to a requirement to collect and share information on universities' teaching practices to force them to pay more attention to the quality of their often poor undergraduate teaching. It should also be tied to a requirement that undergraduates are required to spend more of their time studying. The days of 'an essay or two per term' should end.

Many of those now attending university courses in the UK and USA are wasting their time, and their own and taxpayers' money, and would be better off in jobs or work-based training.[163] In many third-rate HE institutions, there is a large amount of 'social science' work (in economics, anthropology, sociology, literary theory, and so on) of questionable value both from an intellectual perspective and from the perspective of the students' job prospects.[164] Reform of the long tail of HE and FE is crucial and the imminent 'tsunami' of MOOCs will push the system towards generally useful, occasionally damaging, reform.

---

[161] E.g. the [new joint venture between Harvard and MIT](#), Andrew Ng's [Coursera](#), Sebastian Thrun's [Udacity](#) project, [Udemy](#). Thrun explains [here](#) how in 2011 he put his Stanford AI course free on the web and attracted hundreds of thousands of students. Inspired by the feedback, he has left Stanford to start a new free, online educational project. *Science* and *Nature* are developing their own online multimedia textbooks, e.g. Nature's 'Principles of Biology'. New apps give access to content such as [Leonardo's drawings](#). The One Tablet Per Child Project (Negroponte, MIT) recently airdropped tablets, with educational apps, into two remote Ethiopian villages; no instructions were attached (nobody could have read them) and software logs use, reports how they are used, and allows researchers to update the software (but no direct contact with the children). Peer-to-peer services like [OpenStudy](#) allow teachers to share content. Leading game company, [Valve](#), has launched '[Teach With Portals](#)' to help teachers use the game Portal2 to teach STEM, and the National STEM Video Game Challenge aims to encourage games that teach STEM (see below for 'serious games'). A note of warning on games [here](#).

[162] While some subjects (such as physics) benefit from concentration of resources, others (particularly maths) benefit from diverse funding.

[163] [New data](#) suggests that large numbers of US undergraduates are learning little if anything in college courses. Students reported studying little more than 12 hours per week; few courses required > 40 pages of reading per week or writing as much as 20 pages in a semester. Meanwhile costs have risen much faster than inflation and there is a trillion dollars in outstanding student loan debt.

[164] This is partly because of awful intellectual fads associated with French literary theory (Foucault, Derrida and Lacan). As Steven Pinker wrote recently, 'The humanities have yet to recover from the disaster of postmodernism, with its defiant obscurantism, dogmatic relativism, and suffocating political correctness' (*[Science Is Not Your Enemy, 2013](#)*).

MOOCs and their successors will potentially allow not only the enhancement of existing excellent courses but also may encourage the transition from lectures to empirically-based teaching that Wieman *et al* have advocated. Universities are beginning to tell their students to watch lectures from other universities online while redeploying their own resources to more tutorials and small group discussions, seminars and so on, on the Wieman model. The Georgia Institute of Technology's cut-price experiment with offering an online Master's in Computer Science is being closely watched in HE. Already people are preparing physical test centres where people doing MOOCs can prove their identity and do strictly certified tests. Such ideas will filter to schools and merge with teacher training programmes (see below): already Professor Warner at Cambridge is planning a MOOC element in his new 16-18 Physics project and great schools will surely follow great universities in experimenting with putting classes online. When such courses can be properly assessed with identity fraud eliminated, how long will a National Curriculum and state-controlled exams such as GCSEs survive?[165] The UK Department for Education has already made clear that it is handing control of A Levels back to Universities and exam boards.[166] It should change the school accountability system to encourage the spread of MOOCs and certified testing. This will accelerate rivals to A Levels and GCSEs, in the form of real tests of advanced skills, and the latest National Curriculum will hopefully be the last.

7. *Computer Science and 3D printers: bits and atoms, learning and making…* The old *studia humanitatis* involved the *trivium* (grammar, logic, rhetoric) and *quadrivium* (arithmetic, geometry, astronomy, music). The ability to make things was not included. *We can now reunite making things with the training of the intellect*: 'Instead of trying to interest kids in science as received knowledge, it's possible to equip them to *do* science, giving them both the knowledge and the tools to discover it' (Gershenfeld).[167] The spread of good Computer Science courses in schools and intelligent use of 'personal fabricators' - a mix of programming, design, science, modelling, art and engineering - would bring big benefits, as MIT and some schools have already shown: 'I've even been taking my twins, now 6, in to use MIT's workshops; they talk about going to MIT to make things they think of rather than going to a toy store to buy what someone else has designed' (Gershenfeld). The English 'ICT' curriculum has been replaced by a 'Computing' curriculum including programming skills, Computer Science GCSE and a new teacher training programme have been introduced, the Design and Technology curriculum has been revised to include 3D printing, and there are pilots underway with personal fabricators. Much needs to be done in converting school practice in these fields. This report summarises many educational projects using 'FAB' technology. (There are also experiments with introducing '*network science*' (connecting maths, physics, and computer science) to teenagers.)

---

[165] These three interesting blogs on MOOCs (1,2,3). Arguably, only research requires something like existing Universities. Similar approaches are being applied to Universities' administration. There are also experiments with analysis of students' facial expressions to discern what they are confused about, and proposed experiments to build these into MOOCs (using webcams) so that software will study when students are baffled by the MOOC's material, potentially allowing customisation.

[166] Despite constant complaints about A Levels and 'political interference', universities resisted taking back responsibility for A Levels partly because they (institutionally) are more concerned to avoid political accusations of elitism than they are to improve the system. Many individual academics have been keen to help reform A levels but their universities generally discourage this. Until universities regain responsibility for A Levels, they will remain poor. The interaction of the DfE, exam boards, learned societies, and universities is dysfunctional and should be dismantled.

[167] The modern conservative philosopher, Oakeshott, opposed what he regarded as the rationalist Left's exclusion of practical skills from education: '[W]hat is true of cookery, of painting, of natural science and of religion is no less true of politics: the knowledge involved in political activity is both technical and practical … as in all arts which have men as their plastic material… Technical knowledge can be learned from a book… Practical knowledge can neither be taught nor learned, but only imparted and acquired. It exists only in practice and the only way to acquire it is by apprenticeship to a master - not because the master can teach it (he cannot), but because it can be acquired only by continuous contact with one who is perpetually practising it… Rationalism is the assertion that what I have called practical knowledge is not knowledge at all, the assertion that, properly speaking, there is no knowledge which is not technical knowledge.'

One of the most important discoveries of the 20[th] Century was the discovery of profound connections between a) theories of logic and computation and b) the physical world - the ability of 'Universal Turing Machines' to simulate physical systems, and the way that physical systems compute information (cf. Section 3 and Endnote on Computation). Connections between these subjects are not explored in traditional school curricula but the arrival of Computer Science and 3D printers will, hopefully, change this.[168]

8. *Teacher hiring, firing, and training…* Much debate around teacher training involves (implicitly or explicitly) the idea of creating tens of thousands of 'excellent' teachers. However, only a small proportion of the small number of 'significantly more than averagely talented people' are, without huge cultural changes, going to want to be teachers (at least in the conventional career sense). Finland and Singapore recruit teachers from among the top ~¼ of graduates, train them thoroughly, weed out those unsuitable, and they come at or near the top of international measures but others have generally failed to mimic this successfully and studies have failed to identify particular characteristics of a 'good teacher'. It is therefore hard for regulation to produce desired results; we do not know what set of regulations to write in order to generate 'great teachers', though recruiting from among the best rather than worst educated seems a sensible start. However, there is evidence that heads reach consistent judgements within broad ranges.[169] The conclusions seem to be: states do not understand (notwithstanding outliers) what defines good teachers well enough to justify detailed centralised control of training (and the odds of replicating the Finland system in England or America *via* state action seems politically impossible for the moment); we should give heads power to remove who they consider to be poor teachers; we should let schools experiment with rigorous methods of judging teacher effectiveness to improve our knowledge. Doug Lemov's work seems to be good evidence for a strong focus on making teacher training empirical and therefore classroom-based, using immediate feedback on what works (Lemov influenced reforms 2010-13).

In England, the Government has raised the entry requirements for those wishing to enter state-funded training programmes (better degrees required, tougher tests); paid much higher bursaries for trainees with good degrees in maths, physics, chemistry, foreign languages, and computer science (though there remains insufficient focus on hard skills in these areas); removed most of the national pay framework so schools can pay good teachers more and not give automatic pay rises to poor performers; made it easier to hire people from outside the state system to break the grip of Unions and bureaucracy (e.g making Qualified Teacher Status voluntary and allowing schools to

---

[168] In this context, it is interesting that Turing and von Neumann were world-changing mathematicians who also paid extremely close attention to the engineering details of computers; they did not confine themselves to the abstract theory, and indeed this was a bone of contention with the IAS, which abandoned its computer project after von Neumann's death.

[169] Hanushek argues (2011) that rigorous studies have consistently shown that a teacher 1 standard deviation better than average raises student performance by ~0.2-0.3 SDs in a year. Assume a long-term gain of 70% of 0.2 SDs. A +1SD teacher produces an average annual gain of ~$20,000 per child in earnings, so ~$400,000 for a class of 20 (and a low-performer -1SD would reduce earnings by the same amount). US standards could reach Finland's level if the bottom 5-10% of teachers were replaced with average teachers, and this would increase the average growth rate of the USA GDP by about 1% per year. *However*, despite hundreds of studies showing that teacher quality is the most important factor in pupil performance, studies have *not* identified any particular characteristics (masters degrees, experience (after the first 3 years) show almost no effect) and 'value-added' measures based on SATs alone are unreliable. Cf. '*Value-added Measures and the Future of Educational Accountability*', Harris (*Science*, 12/8/2011). This 2012 research (Levitt) found significant benefits from performance-related teacher pay, though the design of the system is unsurprisingly important. The UK's STRB examined global evidence (2012). This paper found the current UK pay system is damaging. This study by the Gates Foundation (2012) suggests that it *is* possible to devise reliable measures of effective teachers that perform robustly on randomised trials.

hire heads from outside education); made it easier to monitor and remove bad teachers;[170] established a network of 'teaching schools' (though this is embryonic and it is hard to know how effective it is); and shifted control of teacher training to schools *via* 'School Direct', so schools can recruit people themselves and choose how and where to train them, instead of training colleges controlling the process.[171] The best Academy chains are also developing their own teacher training programmes using School Direct (there are some early signs of this) and on-the-job training programmes. The DfE should give away all control of training and just monitor assessments.

Perhaps *we should also stop discussing schools as if we are going to have a quarter of a million 'brilliant' teachers and instead think about what to do with tens of thousands of roughly averagely talented people.* While some children will always be blessed by a brilliant teacher, by definition this is not a scaleable solution to our problems: real talent is rare, mediocrity ubiquitous. Whilst heads need to be flexible enough to allow talented people to experiment, we also need schools in chains that spread proven approaches (and '90% solutions') *without* relying on innovation, inspiration, and talent. 'Direct Instruction' (DI), in which teachers follow tightly constrained rules to deliver lessons, is generally ignored in English education debates despite its proven effectiveness in randomised control trials around the world. However, standards might improve substantially if thousands of roughly averagely talented teachers simply used DI instead of reinventing square wheels. It will also be possible for teachers to learn alongside pupils using MOOCs.

Such developments could provide at least a partial escape from traditionally insoluble bottlenecks such as a small number of brilliant teachers mostly rationed for richer families.

9. *Prizes…* We could also use prizes creatively (cf. Endnote), such as a prize of £X million for 'the first programme that reliably gets Y% of pupils who have scored 100 in repeated IQ tests to level Z in maths'. This sort of Grand Challenge, so successful in other fields, could help us break out of current assumptions about the proportions of pupils who can master 'advanced' material and might also encourage more research into precisely what different children are able to learn.

10. *Simplify the appallingly complicated funding system, make data transparent, and give parents real school choice…* In 2010, the current Government inherited a funding system so complicated that no Department for Education official claimed to understand more than a fraction of it, the IT could not cope with it, major administrative errors were endemic, there were (and are) huge funding disparities across the country (some pupils attract double the funding of others), and accidental distortions of incentives were commonplace. Some progress has been made towards simplifying it (2010-13) and the DfE's plan is to introduce in 2015 a '*national funding formula*' based on a flat per pupil amount across the country with as few tweaks as possible. This is inherently worthwhile (fairer, simpler, less bureaucratic) and is also a necessary condition for making it easier for good schools to expand.

Many powerful parts of the system strongly opposed *data transparency*: e.g. the unions were happy with the 'equivalent' scam whereby often worthless courses were treated as 'equivalent' to GCSE

---

[170] Although lots of research shows how important feedback is (and the plummeting cost of putting video online makes it much easier), unions defended silly rules (now gone) that formally limited the time heads could watch teachers. The DfE removed the extra barriers to dismissing bad teachers imposed by the state, so that it is now roughly as easy (legally) to remove a bad teacher as it is to remove someone in the private sector, though there remain considerable non-legal barriers. The EU's 'TUPE' rules also make it hard and therefore expensive for new management to replace teachers when a school is taken over; changing TUPE has been difficult but hopefully BIS will significantly change it soon. In 2011, the education sector accounted for 53% of all working days lost to industrial action (ONS).

[171] One of the teacher training colleges usually described as 'one of the best three' has objected to the Cambridge maths and physics projects on the grounds that they are 'elitist' and we should not try to introduce such advanced ideas into schools.

Physics or Latin, and many oppose opening the National Pupil Database to researchers. To some extent this has been overcome but there is still far to go. League tables have been changed but the problems not eliminated. One of the main problems is that state schools have defined success according to flawed league table systems based on flawed GCSEs while private schools have defined success according to getting pupils into elite universities and therefore taught *beyond* A Levels.[172]

There is no doubt that some 'charter school' chains / Academy chains in the US/UK have done much better than average state schools and better than state schools in the richest areas. The dramatic experience of post-Katrina New Orleans showed that the replacement of many bad state schools (physically destroyed) by charter schools produced large improvements. The KIPP chain ('Knowledge Is Power') is proof that charters can bring significant improvement for poor parents by applying a formula of tough discipline, significantly extended school hours, focus on reading and maths, and better teachers. Continued denials by unions are not sustainable in the face of evidence such as Mathematica's 2013 study which shows large gains for KIPP pupils compared with those who failed to get into KIPP *via* their entry lottery. Large numbers of poorer children can clearly do far better at school than they do.[173]

However, it is also clear that *charter schools / Academies are no panacaea* (only certain regulatory systems and school types produce clearly good results) and the successes of a small number of brilliant organisations are not necessarily scaleable. Across America, some states have managed the charter system badly and others well. The best charter system is Boston's. The 2009 study by Tom Kane (Harvard) and the 2013 Stanford CREDO study both showed very high performance by Boston's charter schools, with children gaining about *an additional year of learning per year* compared with normal state schools (and most other charters). The results were so high that the study's publication was delayed to check the figures.

> 'Compared to the educational gains that charter students would have had in [traditional public schools], the analysis shows on average that students in Boston charter schools have significantly larger learning gains in both reading and mathematics. In fact, the average growth rate of Boston charter students in math and reading is the largest CREDO has seen in any city or state thus far. At the school level, 83 percent of the charter schools have significantly more positive learning gains than their [traditional public school] counterparts in reading and math, while no Boston charter schools have significantly lower learning gains…

> '[T]he typical student in Massachusetts charter schools gains more learning in a year than his [traditional public school] counterparts, amounting to about one and a half months of additional gains in reading and two and a half months in math. The advantage in learning in Boston charter schools equates to more than twelve months of additional learning in reading and thirteen months more progress in math. These outcomes are consistent with the result that charter schools have significantly better results than [traditional public schools] for minority students who are in poverty.' Stanford's 2013 CREDO Study.

It is also the case that many other charter chains have done no better, and sometimes worse, than

[172] Cf. Michael Nielsen op. cit. on data transparency; his work inspired the policy of opening the National Pupil Database to researchers. While there are problems and sometimes absurdities with how the Freedom of Information Act is implemented, and parts of government where its influence should be limited or eliminated, its influence on education policy is generally benign and with tweaks could be more so.

[173] Angrist (MIT) argues that 'We're showing dramatic gains in middle school and later for kids who come in at a very low baseline' in urban charter schools. The idea that children cannot compensate for disadvantages they face early in life is 'a compelling narrative but it's not true.' This Angrist (2013) paper summarises recent findings on the successes of Boston's charter schools.

normal public schools. Similarly, in England while there is clear evidence for the success of chains such as ARK and Harris, there is also no doubt that many Academies are badly run and many, like regular state schools, have gamed the league table system (under huge pressure from Whitehall). While some Free Schools are an inspiration (e.g. Greenwich Free School), some will fail and have predictable disasters from disastrous teaching to financial fraud. *Supporters of charters and Academies need to focus on how to get a regulatory system that deals promptly with failure and allows successful organisations to expand.*

While there will always be private schools and individual 'state' schools, it is vital that school systems make a transition from a cottage industry to one dominated by well-managed chains that benefit from professional management, economies of scale, and spreading empirically demonstrated methods. The English state system could improve enormously *without any innovation other than at the management level*, merely by replacing failed methods and people. Managing schools is much easier than being a brilliant maths teacher and requires only the import of competent (not brilliant) professional managers from outside the education world. (This is easier now because the Government made it possible for people to become heads without the 'NPQH' qualification.)[174]

The other things described above (1-9) could be done even if one disagrees with the idea of a decentralised system driven by parent choice and prefers the old hierarchical system run by MPs, Unions, and civil servants. However, why would one prefer that? Hopefully, recent reforms will push the English system towards one in which the state provides a generous amount of funding per pupil which parents can spend in any school they wish, thus breaking down the barrier of private / state school, while the Department for Education does little more than some regulatory, accountancy, and due diligence functions and has nothing to do with the management of schools, the curriculum, exams, or teacher training.[175]

The widespread realisation that the dominant models for schools and universities are failing offers some hope that change will be fast and effective. However, given that so much success in education depends on motivation, it is also possible that the spread of new methods and technologies will mainly benefit those who are *already* able and motivated and will, therefore, as well as raising average performance, *increase* disparities in outcome. It is generally ignored in education debates that an effective decentralised system that is personalised to get the best out of children with different genetic aptitudes will *increase*, not decrease, gaps in performance - something that some egalitarians foresaw and therefore they opposed 'meritocracy'.[176]

---

[174] Many market-sceptics argue that many parents are feckless and will not pay attention to quality *but competition benefits those who do not pay attention to quality*: e.g. supermarkets have to be sensitive to price competition because enough people, not all, are price sensitive. Many dislike the idea of schools 'going bankrupt' (and think there is something inherently bad about such a thing) but the evidence from post-Katrina New Orleans is that despite extreme disruption children have benefited from bad schools being forced to close; the worst 'disruption' is the lifelong disruption of spending one's whole childhood in a bad school. Cf. Tim Harford blog, op. cit.

[175] Few would support the idea that politicians should 'run' universities (including those academics who like politicians 'running' almost everything else). There is little reason to think that politicians know more about schools than they do about university teaching or research. The main obstacle to the removal of political control is 'chimp politics': claiming 'we will do X for state schools' is a way to signal priorities to voters and many voters would vote for someone who says 'I will do X for state schools therefore it is obvious that I care more about them than Y who wants to let profit-making companies run schools in a dog-eat-dog system.'

[176] E.g. Michael Young, author of *The Rise of the Meritocracy*. Media discussion about wealth, ability and 'social mobility' is explored in this paper by Civitas. The 'Feinstein graph' purporting to show bright poor pupils overtaken by less intelligent rich pupils seems to have been the statistical phenomenon 'reversion to the mean': 'There is currently an overwhelming view amongst academics and policymakers that highly able children from poor homes get overtaken by their affluent (but less able) peers before the end of primary school... [T]he methodology used to reach this conclusion is seriously flawed' ('*The use and misuse of statistics in understanding social mobility*', Jerrim and Vignoles 2011).

Training and Decisions

'Politics is showbiz for ugly people.' A Los Angeles proverb.

'No matter who you are, most of the smartest people work for someone else.' Bill Joy, founder of Sun Microsystems, creator of Java programming language.

How should we 'think about thinking'? Why is effective action so hard and why do most endeavours fail? (Cf. Endnote on Failure.) Given we can all read Thucydides and Sun Tzu and there are many outstanding studies of the dynamics of successful and failing organisations, why do we nevertheless keep making the same sorts of mistake, with cycles of senior people being presented with 'lessons learned' reports on disasters only to repeat the same failings? Why do so many with high IQs leave the world's best universities and immediately start doing things that we know will lead to predictable problems and failure? How do human decisions shape organisations, institutions, and cultures?

Humans have made transitions from numerology to mathematics, from astrology to astronomy, from alchemy to chemistry, from witchcraft to neuroscience, from tallies to quantum computation. However, the vast progress made in so many fields is clearly not matched in standards of government where it remains fiendishly difficult to make progress in a transition from (i) largely incompetent political decision-makers making the same sort of mistake repeatedly and wasting vast resources while trying to 'manage' things they cannot, and should not try to, 'manage', to (ii) largely competent political decision-makers who embed some simple lessons, grasp what it is reasonable to attempt to 'manage' and have the ability to do it reasonably well while devolving other things and adapting fast to inevitable errors.

Current curricula (in even elite schools and Universities) do not introduce people to material that is vital if one is to attempt crude, trans-disciplinary integrative thinking about the problems in this paper. Many senior decision-makers did degrees such as English, History, and PPE. They operate with fragments of philosophy, little maths or science, little knowledge of evolutionary systems (biological or cultural), and little understanding of technology. They are sometimes relatively high scorers in tests of verbal ability but are rarely high scorers in tests of mathematical ability or have good problem-solving skills in cognitively hard areas. They have to make constant forecasts but have little idea about how to make them, how the statistics and computer models underlying them work, or how to judge the reliability of their own views. A [recent survey](#) of 100 MPs by the Royal Statistical Society found that only 40% of MPs correctly answered a simple probability question (much simpler than the type of problem they routinely opine on): 'what is the probability of getting two heads from flipping a fair coin twice?' Despite their failures on a beginner question, about three-quarters nevertheless said they are confident in their ability to deal with numbers (cf. Endnote on Bayes and risk).

The problem of bad predictions is hardly limited to elected politicians or officials - it extends to the entire landscape of political 'experts'. For example, before the 2000 election, the American Political Science Association's members unanimously predicted a Gore victory. A political scientist, Philip Tetlock, has conducted a remarkable study. He charted political predictions made by supposed 'experts' (e.g will the Soviet Union collapse, will the euro collapse) for fifteen years from 1987 and published them in 2005 ('*Expert Political Judgement*'). He found that overall, 'expert' predictions were about as accurate as monkeys throwing darts at a board. Experts were very overconfident: ~15 percent of events that experts claimed had *no chance* of occurring did happen, and ~25 percent of those that they said they were *sure* would happen did not happen. The more media interviews an expert did, the *less* likely they were to be right. Specific expertise in a

particular field was generally of no value; experts on Canada were about as accurate on the Soviet Union as experts on the Soviet Union.

However, some did better than others and he identified two broad categories which he called '*hedgehogs*' (fans of Big Ideas like Marxism, less likely to admit errors) and '*foxes*' (not fans of Big Ideas, more likely to admit errors and change predictions because of new evidence).[177] *Foxes tended to make better predictions.* They are more self-critical, adaptable, cautious, empirical, and multidisciplinary. Hedgehogs get *worse* as they acquire more credentials while foxes get better with experience. The former distort facts to suit their theories; the latter adjust theories to account for new facts. Tetlock believes that the media values characteristics (Big Ideas, aggressive confidence, tenacity in combat and so on) that are the opposite of those prized in science (updating in response to new data, admitting errors, tenacity in pursuing the truth and so on); this means that 'hedgehog' qualities are more in demand than 'fox' qualities, so the political/media market encourages qualities that make duff predictions more likely.[178] 'If I had to bet on the best long-term predictor of good judgement among the observers in this book, it would be their commitment – their soul-searching Socratic commitment – to thinking about how they think' (Tetlock).[179] The extremely low quality of political forecasting is what allowed an outside like Nate Silver to transform the field by applying well-known basic maths.[180] (See below for Tetlock's new project.)

A brief look at economic forecasting, important to politics, demonstrates some of these problems. In the 1961 edition of his textbook, one of the 20[th] Century's most respected economists, Paul Samuelson, predicted that respective growth rates in America and the Soviet Union meant the latter would overtake the USA between 1984-1997. By 1980, he had delayed the date to be in 2002-2012. Even in 1989, he wrote, 'The Soviet economy is proof that, contrary to what many skeptics had earlier believed, a socialist command economy can function and even thrive.'

---
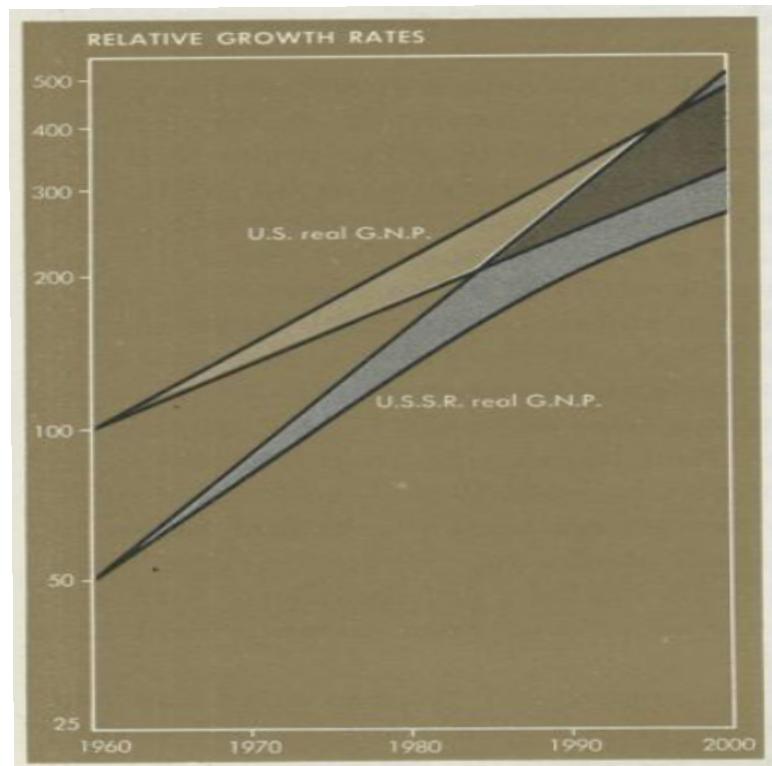
[177] 'The fox knows many little things, but the hedgehog knows one big thing.' Archilochus.

[178] 'There are some academics who are quite content to be relatively anonymous. But there are other people who aspire to be public intellectuals, to be pretty bold and to attach non-negligible probabilities to fairly dramatic change. That's much more likely to bring you attention.' Tetlock

[179] It would be beneficial for someone to create a simple online system to track systematically clear predictions made by UK politicians and political commentators.

[180] 'No official estimates even mentioned that the collapse of Communism was a distinct possibility until the coup of 1989' (National Security Agency, '*Dealing With the Future*', declassified report). The MoD thinks that 'it is still possible to assert, with a fair degree of confidence, what the dominant threat drivers are likely to be out to 2029.' (Really?) Supposedly 'scientific' advice to politicians can also be very overconfident. E.g. A meta-study of 63 studies of the costs of various energy technologies reveals: 'The discrepancies between equally authoritative, peer-reviewed studies span many orders of magnitude, and the overlapping uncertainty ranges can support almost any ranking order of technologies, justifying almost any policy decision as science based' (Stirling, *Nature*, 12/2010).

Chart: Samuelson's prediction for the Soviet economy
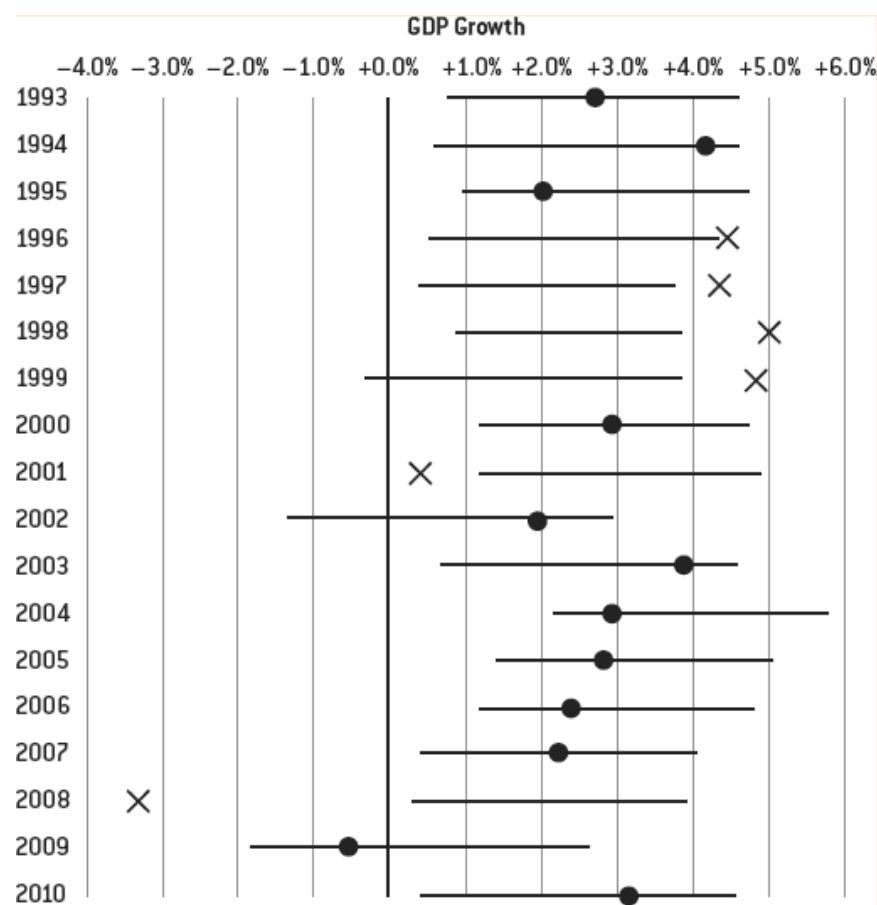
The recent financial crisis also demonstrated many failed predictions. Some, particularly in the world of physics and maths such as Steve Hsu[181] and Eric Weinstein, published clear explanations of the extreme dangers in the financial markets and parallels with previous crashes such as Japan's. However, they were almost totally ignored by politicians, officials, central banks and so on. Many of those involved were delusional. Perhaps most famously, Joe Cassano of AIG Financial said in a conference call (8/2007): 'It's hard for us - without being flippant - to even see a scenario within any kind of realm of reason that would see us losing one dollar in any of those transactions… We see no issues at all emerging.'

Nate Silver recently summarised some of the arguments over the crash and its aftermath. In December 2007, economists in the *Wall Street Journal* forecasting panel predicted only a 38 percent chance of recession in 2008. *The Survey of Professional Forecasters* is a survey of economists' predictions done by the Federal Reserve Bank that includes uncertainty measurements. In November 2007, the Survey showed a net prediction by economists that the economy would grow by 2.4% in 2008, with a less than 3% chance of any recession and a 1-in-500 chance of it shrinking by more than 2%.

---

[181] I read blogs by physicist Steve Hsu from 2005 that were prescient about the sort of collapse that came with the 'quant meltdown' of August 2007 and the crash of September 2008, though the issues were so technical I could not assess them usefully. Almost nobody in Westminster who I emailed them to paid any attention (Alistair Heath is an exception) and many then gave speeches saying 'nobody saw this coming'. I do not mean by this 'I predicted the crash'; I mean it was easier to see what actually happened *as a reasonable possibility* if one paid attention to dissident views, but this is hard to systematise and it is hard to disentangle intelligent dissidents from plausible cranks.

Chart: the 90% 'prediction intervals' for the Survey of Professional Forecasters net forecast of GDP growth 1993-2010

If the economists' predictions were accurate, the 90% prediction interval should be right nine years out of ten, and 18 out of 20. Instead, the actual growth was outside the 90% prediction interval six times out of 18, often by a lot. (The record back to 1968 is worse.) The data would later reveal that *the economy was already in recession in the last quarter of 2007 and, of course, the '1-in-500' event of the economy shrinking by more than 2% is exactly what happened.*[182]

Although the total volume of home sales in 2007 was only ~$2 trillion, Wall Street's total volume of trades in mortgage-backed securities was ~$80 trillion because of the creation of 'derivative' financial instruments. Most people did not understand 1) how likely a house price fall was, 2) how risky mortgage-backed securities were, 3) how widespread leverage could turn a US housing crash into a major financial crash, and 4) how deep the effects of a major financial crash were likely to be. 'The actual default rates for CDOs were more than two hundred times higher than S&P had predicted' (Silver). In the name of 'transparency', S&P provided the issuers with copies of their ratings software allowing CDO issuers to experiment on how much junk they could add without losing a AAA rating. S&P even modelled a potential housing crash of 20% in 2005 and concluded its highly rated securities could 'weather a housing downturn without suffering a credit rating downgrade.'

---

[182] There have been eleven recessions since 1945 but people track millions of statistics. Inevitably, people will 'overfit' many of these statistics to model historical recessions then 'predict' future ones. A famous example is the Superbowl factor. For 28 years out of 31, the winner of the Superbowl correctly 'predicted' whether the stock exchange rose or fell. A standard statistical test 'would have implied that there was only about a 1-in-4,700,000 possibility that the relationship had emerged from chance alone.' Just as *someone* will win the lottery, some arbitrary statistics will correlate with the thing you are trying to predict *just by chance*. (Silver)

Unsurprisingly, Government unemployment forecasts were also wrong. Historically, the uncertainty in an unemployment rate forecast made during a recession had been about plus or minus 2 percent but Obama's team, and economists in general, ignored this record and made much more specific predictions. In January 2009, Obama's team argued for a large stimulus and said that, without it, unemployment, which had been 7.3% in December 2008, would peak at ~9% in early 2010, but with the stimulus it would never rise above 8% and would fall from summer 2009. However, the unemployment numbers after the stimulus was passed proved to be even worse than the 'no stimulus' prediction. Similarly, the UK Treasury's forecasts about growth, debt, and unemployment from 2007 were horribly wrong but that has not stopped it making the same sort of forecasts.[183]

Paul Krugman concluded from this episode: *the stimulus was too small*. Others concluded it had been a waste of money. Academic studies vary widely in predicting the 'return' from each $1 of stimulus. Since economists cannot even accurately predict a recession *when the economy is already in recession*, it seems unlikely that there will be academic consensus soon on such issues. Economics often seems like a sort of voodoo for those in power - spurious precision and delusions that there are sound mathematical foundations for the subject without a proper understanding of the conditions under which mathematics can help (cf. Von Neumann's explanation in Section 1).[184]

Kahneman summarises decades of research on expertise:

> '… the evidence from more than fifty years of research is conclusive: for a large majority of fund managers, the selection of stocks is more like rolling dice than like playing poker. Typically at least two out of every three mutual funds underperform the overall market in any given year. More important, the year-to-year correlation between the outcomes of mutual funds is very small, barely higher than zero. The successful funds in any given year are mostly lucky; they have a good roll the dice...

> 'To know whether you can trust a particular intuitive judgment, there are two questions you should ask: *Is the environment in which the judgment is made sufficiently regular to enable predictions from the available evidence?* The answer is yes for diagnosticians, no for stock pickers. *Do the professionals have an adequate opportunity to learn the cues and the regularities?* The answer here depends on the professionals' experience and on the quality and speed with which they discover their mistakes. Anesthesiologists have a better chance to develop intuitions than radiologists do. Many of the professionals we encounter easily pass both tests, and their off-the-cuff judgments deserve to be taken seriously. In general, however, you should not take assertive and confident people at their own evaluation unless you have independent reason to believe that they know what they are talking about.' (Emphasis added)

---

[183] Many of these wrong forecasts were because the events were '*out of sample*'. What does this mean? Imagine you've taken thousands of car journeys and never had a crash. You want to make a prediction about your next journey. However, in the past you have never driven drunk. This time you are drunk. Your prediction is therefore *out of sample*. Predictions of US housing data were based on past data but there was no example of such huge leveraged price rises in the historical data. Forecasters who looked at Japan's experience in the 1980's better realised the danger. (Silver)

[184] I have found people in politics with training in economics generally (some exceptions) to be overconfident about their understanding and predictions; too trusting that economics is 'objective' and 'scientific' while unaware of the low esteem in which their subject is held by natural scientists; naively reductive; and quick to change predictions whilst maintaining 'I've always thought…' I suspect economics training may encourage a spurious self-confidence in flawed quantitative models and to the extent that PPE has spread conventional economic thinking and forecasting models it has damaged political decision-making. In contrast, natural scientists with whom I have discussed politics have seemed both cleverer and much more aware of their own ignorance. 'I importune students to read narrowly within economics, but widely in science...The economic literature is not the best place to find new inspiration beyond these traditional technical methods of modelling' (Vernon Smith).

As well as poor judgement and forecasting skills, political 'experts' often have little training in, experience of, or temperament for managing complex processes or organisations. Many of those in senior positions have *never* worked in a non-dysfunctional entity - they have no idea of what it is like to work in an effective operation (while people in other systems often experience a contrast between places with an unusual culture of excellence, normal mediocrity, and dysfunction). They often have little feel for how decisions will ripple through systems (including bureaucracies) into the real world.[185] The fact that they often cannot manage the most simple things (like answering correspondence) rarely undermines their confidence that they should have power over the most important things. Their approach to management often involves attempts at detailed micromanagement and instructions rather than trusting to *Auftragstaktik* (give people a goal and let them work out the means, do not issue detailed instructions), which requires good training of junior people (itself rare).

On one hand, many take pride in *not* having a plan.[186] On the other hand, politics is dominated by discussion of 'strategy' and 'priorities', but few know how to think strategically (or even what 'strategy' is) or how to make and stick to priorities (see below) and 'planning' often does not extend beyond about ten days. 'Strategy' is much mentioned but little studied. *Strategy* is not 'goals', 'vision' or rhetoric. *Strategy focuses action on crucial problems to connect operations to aims*; it requires diagnosis, a guiding policy, and coherent action.[187] Good strategy requires choices, choices require *not* doing some things, and some people will be upset at not being 'a priority'; therefore, good strategy is by definition hard for politicians to articulate even if they can develop it.[188]

---

[185] 'The man of system, on the contrary, is apt to be very wise in his own conceit; and is often so enamoured with the supposed beauty of his own ideal plan of government, that he cannot suffer the smallest deviation from any part of it. He goes on to establish it completely and in all its parts, without any regard either to the great interests, or to the strong prejudices which may oppose it. He seems to imagine that he can arrange the different members of a great society with as much ease as the hand arranges the different pieces upon a chess-board. He does not consider that the pieces upon the chess-board have no other principle of motion besides that which the hand impresses upon them; but that, in the great chess-board of human society, every single piece has a principle of motion of its own, altogether different from that which the legislature might choose to impress upon it.' Adam Smith.

[186] Many think that Macmillan's 'events, dear boy, events' is a maxim to be emulated. 'Above all, no programme' (Disraeli). 'I distrust anyone who foresees consequences and advocates remedies for the avoidance of them' (Halifax before the war). When it was suggested to Cadogan in 1944 that the Foreign Office have a two-year policy statement updated six-monthly, he answered, 'That way lies bedlam.' (I have sympathy with the last, but the point stands.)

[187] Clausewitz described military strategy as 'the use of the engagement for the purpose of the war' and says the strategist 'must therefore define an aim for the entire operational side of the war that will be in accordance with its purpose.' Colin Gray defines military strategy as 'the direction and use that is made of force and the threat of force for the ends of policy'. The first use of 'strategy' in a sense beyond narrow generalship was in 1777 in French and German and prior to 1810 English dictionaries did not contain a 'strategy' entry. 'Strategy was not recognized linguistically as a function distinctive from statecraft or generalship prior to the late 18th century. Polities did not have permanent or even temporary schools and military staff charged with "strategic" duties. Policy and strategy, though logically separable, usually were all but collapsed one into the other.' (Gray, *Schools for Strategy*, 2009).

[188] Bad strategy is identified by: fluff (vague, grandiloquent rhetoric), ignoring important problems, mistaking goals for strategy, and setting bad (or contradictory) 'strategic objectives'. It is not miscalculation. It is sometimes a substitution of belief for thought. Now it is often produced *via* a powerpoint template, with visions, mission statements, core values, strategic goals, lists of initiatives etc - all avoiding the hard questions (Rumelt, 2011). Many significant political problems that this author has discussed with participants have been characterised by failure at the most basic level: failure to ask, never mind to think clearly with discipline about, questions like - what is the aim?, what are the main problems?, what is the strategy?, how should we judge progress?, who is responsible for what and on what timetable? (E.g. On the euro campaign 1998-2003, most confused the *aim* (keep the pound) with the *strategy* (destroy the opposing alliance so Blair could not hold a referendum), just as now many confuse aims ('leave the EU' or 'win a referendum'?) Rarely does one hear people ask 'how have intelligent people tried to solve this problem already and what happened?'

Misunderstanding of strategy, and the proliferation of rhetoric masquerading as strategy, causes huge problems, including with national leaderships' attempts to define 'national strategy'.[189]

Many are dominated by the pursuit of prestige, refusal to admit errors,[190] and by a desire to react to media coverage, which is much simpler than real problem-solving or management and provides a comforting illusion of activity.[191] They are often persuasive in meetings (with combinations of verbal ability, psychological cunning, and 'chimp politics') and can form gangs. Many 'political' institutions select *for* those who pursue prestige and suppress honesty and *against* those with high IQs, a rational approach to problem-solving, honesty and selflessness; *they are not trying to recruit the best problem-solvers*. A refusal to face one's flaws can be a perverse 'asset' in many political institutions. The skills of research scientists are hard to harness in political institutions (consider how easy it is for a) a lawyer and b) a cutting-edge scientist to have a senior political job without sacrificing their career) and Feynman's ideal of 'a kind of scientific integrity, a principle of scientific thought that corresponds to a kind of utter honesty - a kind of leaning over backwards' - this is practically the opposite of how politics works (inevitably to some extent, but much of the trickery is simply *fooling yourself*).[192] This blog by physicist Steve Hsu, *Creators and Rulers*, discusses the differences between genuinely intelligent and talented 'creators' (e.g. scholars, tech start-ups) and the 'ruler' type that dominates politics and business organisations (e.g. CEOs with a history in law or sales). This seems accurate to this author and it is worth considering that the 'Ruler' described

---

[189] Cf. these two papers (1, 2) on the problems with US strategic planning. The generally accepted best recent example of strategic planning in the US national security establishment was Eisenhower's Project Solarium in 1953. Unusually, Eisenhower valued dissent and argument and did not encourage consensus. 'I have been forced to make decisions, many of them of a critical character, for a good many years, and I know of only one way in which you can be sure you have done your best to make a wise decision … that is to get all of the [responsible policymakers] with their different points of view in front of you, and listen to them debate.' It seems that the classified thinking on problems of nuclear deterrence during the Cold War was often terrifyingly misguided (cf. Payne) and prone to classic misunderstandings that could easily have sparked catastrophe. These problems are worse in a world of garage biology, wider proliferation, faster communication, greater opportunities for psychological manipulation etc. Some of the problems regarding intelligence analysis are discussed in *Psychology of Intelligence Analysis*, by Richards Heuer. Arguably, Britain has badly mismanaged the strategic analysis of most major issues (e.g. the rise of Germany, the EU, technology) since the famous dinner party in the summer of 1862 when Bismarck, briefly visiting London before his fateful appointment as Prime Minister of Prussia, discussed the European scene with our leading politicians and wrote to his wife in sarcastic terms, 'It was very pleasant there but the British ministers know less about Prussia than about Japan and Mongolia.'

[190] 'Strength of character can degenerate into obstinacy. The line between the two is often hard to draw in a specific case ... Obstinacy is not an intellectual defect; it comes from reluctance to admit that one is wrong... Obstinacy is a fault of temperament' (Clausewitz). Doing hard things requires will power, as well as superior understanding of reality, in order to maintain motivation in the face of adversity, but such will power can easily become catastrophic blindness to reality and a refusal to admit error. Since it is usually impossible to know when will power might, if maintained, lead to the collapse of others' will and therefore provide new opportunities ('never give up'), it is generally impossible to decide this problem 'rationally'. E.g Steve Jobs partly grasped the reality of some dynamics better than others, but he also had what people who knew him called 'a reality distortion field': this sometimes led to disaster but was *also* credited with getting the best out of people and managing to do things others thought impossible. As one of his colleagues said, 'His reality distortion is when he has an illogical vision of the future… You realize that it can't be true, but he somehow makes it true.' 'To close your ears to even the best counter-argument once the decision has been taken: sign of a strong character. *Thus an occasional will to stupidity*' (Nietzsche). E.g Alexander crossing the Makran?

[191] Most 'communication problems' are really failures of management and cannot be solved by 'communication'. ~100% of 'political communication' in Britain is white noise (people are busy and rationally ignore politicians) and the tiny amount that affects the public is often accidental. Few phrases are more common than 'we need a campaign to…' but few things are rarer than a professional campaign that changes millions of people's opinions or feelings. So-called 'strategic communication' is not very hard but is rarely attempted, never mind done, partly because it requires a lot of hard thinking, focus, priorities, facing weaknesses etc - i.e. many things that are psychologically difficult. Most of what people call 'strategic communication' is really just answering phone calls from journalists. In crises, almost everyone panics and spins stories about 'strategy' to journalists whilst its practice dissolves if it ever existed (unlikely). The subject is widely discussed in defence and intelligence circles but also rarely well executed. E.g. The Pentagon knows that the huge amount of effort it has put into 'information operations' did not work in Afghanistan and Iraq (Report).

[192] '[T]he idea is to try to give *all of* the information to help others judge the value of your contribution… The easiest way to explain this idea is to contrast it with advertising' (Feynman, 1974).

there represents with few exceptions the *best* end of those in politics, many of whom are far below the performance level of a successful 'political' CEO.[193] Also, it is possible that those who climb to the top of the hierarchy are more likely to focus only on their own interests.[194]

Most of our politics is still conducted with the morality and the language of the simple primitive hunter-gatherer tribe: 'which chief shall we shout for to solve our problems?' Our 'chimp politics' has an evolutionary logic: our powerful evolved instinct to *conform* to a group view is a flip-side of our evolved in-group *solidarity* and *hostility* to out-groups (and keeping in with the chief could lead to many payoffs, while making enemies could lead to death, so going along with leaders' plans was incentivised). This partly explains the persistent popularity of collectivist policies (see Section 7) and why 'groupthink' is a recurring disaster. Although discussion can improve many types of decision and the 'wisdom of crowds' can aggregate superior judgements by averaging many guesses, many studies show how group discussion of political issues often leads to *worse* decisions than individuals, with less information than was available to the group, would have made. High status people tend to dominate discussion and common information is over-discussed while information unique to an individual, especially a lower status individual, is routinely ignored. The wisdom of crowds only works if many *independent* judgements are aggregated; if social influence distorts the process, one gets disastrous mobs - not the wisdom of crowds. However, despite centuries of discussion and analysis of groupthink disasters, it recurs constantly and *few learn from disasters*, including those who give speeches saying 'we must learn from this disaster'. Students can study many examples in which the same problems lead to the same sorts of errors and failure, particularly the suppression of bad news and a failure to face reality (cf. Endnote on Cognitive Biases).[195]

---

[193] ~40-45% or more of US Fortune 500 CEOs, billionaires, federal judges and Senators attended elite universities whose median standardized test scores are above ~99th percentile for the overall US population: i.e ~1% of the population make up ~50% of the elite group running the country (Wai, 2013). However, even *within* this 1% there are huge differences between the brains and character of a Zuckerberg and an average senator. Cf. Endnote.

[194] Studies such as '*Higher Social Class Predicts Increased Unethical Behavior*' claim that 'the rich are way more likely to prioritize their own self-interests above the interests of other people' (Piff) and even just thinking about money makes people more self-centred. Not only are richer people healthier (less likely to have heart attacks or suffer mood disorders), but they also produce less cortisol (suggesting lower stress levels; cf. studies suggest those at the top of hierarchies suffer *less* stress because they feel a greater sense of control), they are less attentive to pedestrians when driving, and less compassionate when watching videos of children suffering with cancer. There is some evidence that richer parents are more likely to stress individuality while poorer parents stress community interests, and that richer people are more inclined to blame individuals for things going wrong while poorer people blame luck. This article touches on these studies though it should be remembered that many studies of such things are not replicated (Endnote: Problems with science). Further, one of the handful of the most important studies on IQ, personality and scientific and financial success, the 'Terman' study, also shows a *negative* correlation between earnings and agreeableness, and positive correlations between conscientiousness and earnings, and IQ and earnings.

[195] E.g. David Galula's fascinating account of his successful counter-insurgency, '*Pacification in Algeria 1956-8*', discussed how hard it was for armies to remember ancient and modern lessons in this field - and was itself promptly forgotten not only by the Americans (who commissioned it) in Vietnam but for the next forty years until 9/11. McMaster wrote a study of LBJ's failures in Vietnam ('*Dereliction of Duty*'); the suppression of bad news was central. McMaster fought in Iraq in 2003 and saw for himself similar errors repeated. He tried new tactics (small bases embedded in, and helping, the population). He was repeatedly passed over for promotion as superiors suppressed bad news... The reluctance of the NASA bureaucracy to face facts *viz* the Challenger disaster (1986), the 'PowerPoint festival of bureaucratic hyperrationalism', and Feynman's famous pursuit of the facts and exposure of groupthink (which brought the comment from the head of the investigation that 'Feynman is becoming a real pain'), was followed by the Columbia disaster (2003) and another report showing NASA had not learned lessons from the previous disaster, and internal pressure to conform meant 'it is difficult for minority and dissenting opinions to percolate up through the agency's hierarchy'. Political disasters are rarely analysed carefully. E.g. Many doubted that the euro's institutions would work (e.g. Feldstein 1 and 2 and even the ECB's own Otmar Issing). European elites not only rejected such warnings but treated them as the views of the idiotic or malign, and such has been the emotional commitment (cf. Habermas' '*Why Europe Needs a Constitution*') that it is still hard for those elites to consider the euro's / EU's problems rationally.

Robert Trivers, one of the most influential evolutionary thinkers of the last fifty years, has described how evolutionary dynamics can favour not just deception but *self-deception*: conflict for resources is ubiquitous; deception helps win; a classic evolutionary 'arms race' encourages both deception detection and ever-better deception; *perhaps humans evolved to deceive themselves because this fools others' detection systems* (for example, self-deception suppresses normal clues we display when lying).[196] This is one reason why most people consistently rate themselves as above average.[197] Such instincts, which evolved in relatively simple prehistoric environments involving relatively small numbers of known people and relatively simple problems (like a few dozen enemies a few miles away), cause disaster when the problem is 'how to approach an astronomically complex system such as health provision for millions.'

Further, bureaucracies often operate as if *designed* to avoid senior decision-makers confronting reality and discussing hard choices. It is convenient for many to treat senior ministers as if they have the ego and temperament of movie stars who must be protected from brutal reality in the form of unpleasant discussions that may require disagreement, and unfortunately there is a perverse underlying logic to it given people's actual nature.

These problems and dynamics are well known, many have been written about for thousands of years, yet they recur. In response to criticism of their management, politicians often reply, 'politics is always chaos, it's much harder than managing a company, that's why businessmen are so rarely a success in politics.' ('Political' businessmen (the lawyers and salesmen who are often CEOs of public companies) often are a disaster, and entrepreneurial businessmen want to change things

---

[196] Children display deception when just months old (e.g. fake crying). There is 'clear evidence that natural variation in intelligence is positively correlated with deception.' 'We seek out information and then act to destroy it… Together our sensory systems are organized to give us a detailed and accurate view of reality, exactly as we would expect if truth about the outside world helps us to navigate it more effectively. But once this information arrives in our brains, it is often distorted and biased to our conscious minds. We deny the truth to ourselves … We repress painful memories, create completely false ones, rationalize immoral behavior, act repeatedly to boost positive self-opinion, and show a suite of ego-defense mechanisms' (Trivers). Roberta Wohlstetter wrote in '*Slow Pearl Harbors*' re ignoring threats, 'Not to be deceived was uncomfortable. Self-deception, if not actually pleasurable, at least can avoid such discomforts.'

[197] 'But for self-deception, you can hardly beat academics. In one survey, 94 percent placed themselves in the top half of their profession' (Trivers). 'Academics, like teenagers, sometimes don't have any sense regarding the degree to which they are conformists' (Bouchard, *Science* 3/7/09). Even physical scientists who know that teleological explanations are false can revert to them under time pressure, suggesting that such ideas are hardwired and are masked, not replaced, by specific training. 'I don't know how Nixon won, no one I know voted for him' (Pauline Kael, famous movie critic, responding to news of Nixon's 1972 landslide victory). 'The basic mechanism explaining the success of Ponzi schemes is the tendency of humans to model their actions, especially when dealing with matters they don't fully understand, on the behavior of other humans.' Psychiatry Professor Stephen Greenspan, *The Annals of Gullibility* (2008), which he wrote just before he lost more than half his retirement investments in Madoff's ponzi.

quickly so are seen as a menace.) They often do not know any better so substantial change is rarely discussed.[198] This partly explains why so much is mismanaged.[199]

Seldom does a political institution escape these dynamics and demonstrate high performance. Such incidents (sometimes terrifying) are generally ascribed to the happenstance of a particular individual often described as a (evil) 'genius'; once they vanish institutions tend to revert to type and little is learned. No training system that routinely produces politicians who avoid these problems is known (though see below). Similarly, little has been done to develop, patiently and systematically, reliable knowledge about different policies that can be built on, in the way scientists build reliable knowledge-systems: similar reports are re-written, spun, and forgotten, and new gimmicks are routinely drafted hurriedly (and briefed to the media) without learning from similar expensive prior experiments and without anybody asking 'who has already tried to solve this and how did it work out?' While there are some brilliant journalists, and scrutiny often is the only thing

---

[198] These problems have worsened as lawyers and PR people spread, those who successfully manage organisations have no clear route to senior executive government roles, more of the cleverest go into finance rather than public service, and bureaucracies have become more complex while basic administrative skills have declined. In Government, there is constant panic but little urgency: people assume epic timescales (almost everything takes months if not years longer than necessary), vanish early, go on holiday just as the project they're responsible for is launched, switch jobs at the wrong time, leave juniors to deal with emergencies and so on. Many organisations cannot fill senior posts with people who can reliably spell and manage simple processes like answering correspondence. Ministers are a) generally unsuitable to manage departments themselves (even if they were allowed to which they are not) and b) not allowed to select those who do manage them. (Parliaments generally do not provide the right sort of talent pool. 'Particularly notable is the growth in the number of MPs who come to Westminster already with a political background. In 1979 3% of MPs ... were previously politicians/political organisers, compared to 14% in 2010,' House of Commons Library, 2010). Senior officials are occasionally brilliant but are selected by (apparently 'open', actually opaque) dysfunctional processes that often reward 'chimp politics' skills such as flattery. Almost nobody is fired regardless of disaster. Responsibility is opaque. In the private sector stupid management is routine but shareholders can lose their money; in the public sector, stupid management is routine but it can be prolonged with taxpayers' money and failure is often rewarded with hundreds of thousands of pounds (thousands of officials are paid large sums to leave jobs that should not exist). Some Ministers have resorted to FOI-ing their own departments to get vital information. Overall, a culture of excellence is *unthinkable*. In meetings, it is striking how often the clever and able people are on the other side of the table to the MPs, 'special advisors', and officials who make decisions. Every part of the system blames the others and as they say in Russia, 'everyone's right and everyone's unhappy.' Small attempts to remedy these problems are forbidden as 'illegal', large attempts are regarded as 'extreme'. Some excellent officials also want radical change; they need support.

[199] These general factors are compounded by some specific problems in England. The combination of the growth of public law, judicial review, EU regulation, and the ECHR/HRA - combined with civil service practice - has added cost, complexity, and uncertainty. There is no objective view of 'what the law is' in many circumstances so management decisions are undermined many times per day by advice to do things 'to avoid losing a judicial review' the risks of which are impossible to analyse clearly. Legal advice is offered saying that both doing X and not doing X could be 'illegal' leading to Kafka-esque discussions and pseudo-'fair processes' (like 'consultations') designed only to be evidence in court. Internal legal advice makes discussion of regulatory trade-offs tortuous and wasteful; it is always easier to urge 'caution' and 'we'll lose a JR' is an easy way across Whitehall to delay or block change. (It is not well-understood how in many systems, 'safety features' can create more risks. E.g. The chain of events leading to the Piper Alpha disaster (1988) saw a) rules to stop engineers getting too tired on long shifts, and b) safety devices, contribute to the disaster. Sometimes the safety harness drags you off the cliff. The Six Mile Island disaster saw many safety features go wrong and control panels were so confusing they exacerbated errors under pressure. Visual displays that help rather than hinder management of complex crises are a priority. Ch. 6 of *Adapt*, Tim Harford (2011).)

that forces politicians and officials to focus, overall media coverage encourages many of the bad dynamics described above.[200]

What to do?

Adeimantus: 'Themistocles, in the races the man who starts before the signal is whipped.' Themistocles: 'Those who start too late win no prizes.'

'With forty men, you can shake the world.' Old Mongol proverb.

'Never forget that a section can do in the first hour what by nightfall would require a company, the following day a battalion, and by the end of the week an army corps.' Rommell.

'Politics is a job that can really only be compared with navigation in uncharted waters. One has no idea how the weather or the currents will be or what storms one is in for. In politics, there is the added fact that one is largely dependent on the decisions of others, decisions on which one was counting and which then do not materialise; one's actions are never completely one's own. And if the friends on whose support one is relying change their minds, which is something that one cannot vouch for, the whole plan miscarries... One's enemies one can count on – but one's friends!… Politics … is the capacity to choose in each fleeting moment of the situation that which is least harmful or most opportune… With a gentleman I am always a gentleman and a half, and with a pirate I try to be a pirate and a half.' Bismarck

'Scottie [said Thalberg to Fitzgerald], supposing there's got to be a road through the mountain and … there seems to be a half-dozen possible roads… each one of which, so far as you can determine, is as good as the other. Now suppose you happen to be the top man, there's a point where you don't exercise the faculty of judgement in the ordinary way, but simply the faculty of arbitrary decision. You say, 'Well, I think we will put the road there,' and you trace it with your finger and you know in your secret heart, and no one else knows, that you have no reason for putting the road there rather in several other courses, but you're the only person that knows that you don't know why you're doing it and you've got to stick to that and you've got to pretend that you know and that you did it for specific reasons, even though you're utterly assailed by doubts at times as to the wisdom of your decision, because all these other possible decisions keep echoing in your ear. But when you're planning a new enterprise on a grand scale, the people under you mustn't ever know or guess that you're in doubt, because they've all got to have something to look up to and they mustn't ever dream that you're in doubt about your decision. These things keep occurring.' Irving Thalberg, Hollywood legend, to F Scott Fitzgerald.

---

[200] Most political commentary is driven by a cycle of ephemeral emotional waves and follows the polls; if you're ahead, commentators explain your 'brilliant strategy', if you're behind, they explain why your opponents are brilliant. Events are over-interpreted because journalists do not want to face the idea that they are usually spectators of over-promoted people floundering amid chaos - actions must be intended ('their strategy is…'), farcical reality must be tarted up. When a broadcast journalist tells you 'this is a political story', it usually means that it is unlikely that the person responsible will know much about the underlying issue and it will be treated superficially. Everybody complains that 'the public doesn't listen' but *it is rational to ignore most political coverage because it does not convey information about important issues in ways that are useful or even intelligible.* Use of graphics on TV is so crude that reporters will talk about 'the economy achieving lift off' or 'going down the plughole' and use video of a plane taking off or water draining out of a sink. To some extent the media can be manipulated by disciplined people (e.g using exclusives) but the lack of advertising (and license fee) gives the media much more power here than almost anywhere in the world. Few in politics or journalism will face the true nature of public contempt for the political-bureaucratic-media-big business class (exacerbated by the dreadful rules on public and private sector pay that mean widespread payment for failure) or the way in which most political coverage is ignored; it is so dispiriting that the truth is ignored. Imagine what might happen if one senior editor decided to produce something completely different to what is now normal, if the editor focused on what is *important* rather than what is *new*, if he decided to present data and probability properly (e.g. using 'number needed to treat' for drug stories instead of percentages) and explain policy using creative graphics and stories *based on accurate knowledge of public understanding.* Traditional media are in crisis so dramatic change is possible.

'"Tiger, one day you will come to a fork in the road," he said. "And you're going to have to make a decision about which direction you want to go." He raised his hand and pointed. "If you go that way you can be somebody. You will have to make compromises and you will have to turn your back on your friends. But you will be a member of the club and you will get promoted and you will get good assignments." Then Boyd raised his other hand and pointed another direction. "Or you can go that way and you can do something - something for your country and for your Air Force and for yourself. If you decide you want to do something, you may not get promoted and you may not get the good assignments and you certainly will not be a favorite of your superiors. But you won't have to compromise yourself. You will be true to your friends and to yourself. And your work might make a difference." He paused and stared into the officer's eyes and heart. "To be somebody or to do something." In life there is often a roll call. That's when you will have to make a decision. To be or to do. Which way will you go?" Regular speech by legendary fighter pilot, plane designer, and strategist - Colonel John Boyd.[201]

'We tend to let our many subsidiaries operate on their own, without our supervising and monitoring them to any degree. That means we are sometimes late in spotting management problems and that both operating and capital decisions are occasionally made with which Charlie and I would have disagreed had we been consulted. Most of our managers, however, use the independence we grant them magnificently, rewarding our confidence by maintaining an owner-oriented attitude that is invaluable and too seldom found in huge organizations. We would rather suffer the visible costs of a few bad decisions than incur the many invisible costs that come from decisions made too slowly – or not at all – because of a stifling bureaucracy... We will never allow Berkshire to become some monolith that is overrun with committees, budget presentations and multiple layers of management. Instead, we plan to operate as a collection of separately-managed medium-sized and large businesses, most of whose decision-making occurs at the operating level. Charlie and I will limit ourselves to allocating capital, controlling enterprise risk, choosing managers and setting their compensation' (Buffett who has a HQ of two dozen people).

'When superior intellect and a psychopathic temperament coalesce ..., we have the best possible conditions for the kind of effective genius that gets into the biographical dictionaries' (William James).

'We're lucky [the Unabomber] was a mathematician, not a molecular biologist' (Bill Joy).

Andy Groves described how he and Gordon Moore (of 'Moore's Law') managed, 'after aimless wandering had been going on for almost a year', to escape the usual pressures and abandon their doomed continuation of a failed strategy after Japanese competition had suddenly destroyed Intel's core business:

> 'I was in my office with Intel's chairman and CEO, Gordon Moore, and we were discussing our quandary. Our mood was downbeat. I looked out the window… and I turned back to Gordon and I asked, "If we got kicked out and the board brought in a new CEO, what do you think he would do?" Gordon answered without hesitation, "He would get us out of memories." I stared at him, numb, then said, "Why shouldn't you and I walk out the door, come back and do it ourselves?"'

That is what they did and Intel survived and prospered. Groves described another incident in which he, an isolated leader, failed to spot disastrous dynamics around himself until it was almost too late - the 'Intel Pentium crisis'.

> 'But most CEOs are in the center of a for-titled palace, and news from the outside has to percolate through layers of people from the periphery where the action is. For example, I was one of the last to understand the implications of the Pentium crisis. It took a barrage of

---

relentless criticism to make me realize that something had changed and that we needed to adapt to the new environment. We could change our ways and embrace the fact that we had become a household name and a consumer giant, or we could keep our old ways and not only miss an opportunity to nurture new customer relationships but also suffer damage to our reputation and well-being.

'The lesson is, we all need to expose ourselves to the winds of change. We need to expose ourselves to our customers, both the ones who are staying with us as well as those that we may lose by sticking to the past. We need to expose ourselves to lower-level employees, who, when encouraged, will tell us a lot that we need to know. We must invite comments even from people whose job it is to constantly evaluate and critique us, such as journalists and members of the financial community. As we throw ourselves into raw action, our senses and instincts will rapidly be honed again.'

The self-awareness in these two stories is unusual in business or politics.[202]

Better performance and limiting the damage done by human instincts requires better educated and trained leaders. If people of *understanding*, *talent* and *will* enjoy great *training*, then they face reality and know themselves.[203] They spot shadows of the future, they have more *accurate predictions* and, occasionally, Themistoclean *foresight*,[204] more rapid *error correction*, and more *effective action*. They have '*coup d'oeil*' (understanding at a glance) and *Fingerspitzengefühl* (finger-tip feeling) from aims and strategy to operations and crucial details. They plan obsessively ('chance favours only the prepared mind') and adapt to adversity: 'everybody has a plan until they get hit', and many forget it afterwards. They focus on strategic priorities amid chaos. They see how to 'win without fighting' (Sun Tzu) and how to connect to sources of power. They 'make everything as simple as

---

[202] It perhaps explains why Jack Welch noted that new managers found it much easier to identify the bottom 10% (at GE) than those who had been there three years. It is easier for new people to have the emotional disengagement from previous decisions that is usually required to recognise errors. Groves and Moore unusually managed to force themselves to acquire this distance from their own prior decisions and avoid their organisation dying.

[203] 'If you know the enemy and know yourself, you need not fear the result of a hundred battles' (Sun Tzu). Are opponents 'rational' (they can link goals to ends, calculate priorities) and 'reasonable' (they share goals and ethics with you)? Are they overconfident ('I'm the next Alexander'), deluded (they'll never fight over X so I will do...), dreamers (Hitler cancelled the V2 after a dream), whimsical (the Syrian Defense Minister warned terrorists not to harm Italians because of his crush on Gina Lollobrigida), on drugs (Hitler, Mao). Are they happy to 'die beautifully'?

[204] Themistocles and Pericles were singled out by Thucydides for their foresight: 'For Themistocles was a man whose natural force was unmistakeable; this was the quality for which he was distinguished above all others; from his own native acuteness and without any study before or at the time, he was the ablest judge of the course to be pursued in a sudden emergency, and could best divine what was likely to happen in the remotest future… Nobody could foresee with equal clearness the good or bad which was hidden in the future.' On Pericles: 'During the peace while he was at the head of affairs he ruled with prudence; under his guidance Athens was safe, and reached the height of her greatness in his time. When the war came he showed that here too he had formed an accurate estimate of Athenian power. He survived the beginning of war by two years and six months, and, after his death, his *foresight* was even better appreciated than during his life.' A condition of such foresight is courage to face reality.

possible but no simpler' (Einstein).[205] They learn from how others solved similar problems or failed ('good artists borrow, great artists steal') and they innovate instead of losing within established rules. They are fast,[206] fluid, unpredictable,[207] and seize the initiative: 'If revolution there is to be, let us rather undertake it than undergo it' (Bismarck). They can build teams that suppress Chimp Politics, in which everyone knows who is responsible for what and when, and they push them beyond normal limits.[208] They know how to use intelligence[209] and 'influence operations'.[210]

Such people could use *tools* such as:

- Checklists. Gawande's excellent book *The Checklist Manifesto* explains the different approaches to checklists in fields such as airplanes, building, and medicine. Properly built and used, checklists can be very powerful, partly because they force people to consider extremely carefully how things work (e.g. they reveal hidden things like the importance of nurses feeling able to contradict a senior surgeon).[211] He has also explored why some successful ideas spread quickly while others do not. Surgical *anaesthesia* spread across the world within months after the first demonstration in 1846. After Lister published his reports on the success of *antiseptics* to reduce infection (*Lancet*, 1867), based on Pasteur's observations, progress was glacial; twenty years later, hospitals were still full of blood-spattered doctors who did not even wash their hands. Why the difference? 1. Pain was an *immediate and obvious problem*, in the form of screaming people, while germs were an invisible problem. 2. Anaesthesia helped patients and doctors; antiseptics seemed to help only

---

[205] When Steve Jobs returned to a nearly bankrupt Apple in 1997, he found a once-famous brand now regarded as failing, no strategy, no focus, trying to do far too much, massively complicated product lines, leading in little, and dysfunctional decision-making. He simplified and focused. 'The product lineup was too complicated and the company was bleeding cash. A friend of the family asked me which Apple computer she should buy. She couldn't figure it out and I couldn't give her clear guidance, either' (Jobs). 'After a few weeks Jobs finally had enough. "Stop!" he shouted at one big product strategy session. "This is crazy." He grabbed a magic marker, padded to a whiteboard, and drew a horizontal and vertical line to make a four-squared chart. "Here's what we need," he continued. Atop the two columns he wrote "Consumer" and "Pro"; he labeled the two rows "Desktop" and "Portable." Their job, he said, was to make four great products, one for each quadrant.' Isaacson, p.337. He cut all printers and other peripherals. He cut development engineers and software development. He cut distributors and five of the six national retailers. He moved almost all manufacturing offshore to Taiwan. He cut inventory by 80% and began selling directly over the internet. (Rumelt, p.12, 2011). In various ways the machinery of the British state resembles 1997 Apple and needs someone to do what Jobs did: massive focus and simplification, eliminate whole categories of function, and so on.

[206] As Lenin said, sometimes nothing happens in decades and sometimes decades happen in weeks, so be ready. 'Slim started with his own headquarters. Besides weeding out the unfit, subjecting the remainder to grueling road marches, field problems, and marksmanship training, Slim and his small cadre of veterans worked to make the outfit a true operational headquarters. The emphasis was on mobility. Everything, from furniture to map boards to office supplies, was designed to be packed away in sling-mounted footlockers. These could be transported by truck, plane, boat, mule, or porter. Constant drill and strict adherence to load plans allowed the headquarters to displace in a truly remarkable twenty minutes' (Eden on British General William Slim).

[207] Traditional naval doctrine in 1805 was for two fleets to line up parallel to each other and fire. At Trafalgar, Nelson broke his fleet into two columns and attacked Napoleon's line perpendicularly. He ignored conventional wisdom and took a calculated risk to sow chaos in the belief that his superior speed and training would bring victory amid chaos.

[208] Bright, motivated twenty-year olds focused on putting the opponent under hourly pressure can easily cause chaos for famous and supposedly 'expert' figures. When attacked, one's own side is usually the biggest problem. As Bismarck said, 'One's enemies one can count on – but one's friends!'

[209] 'A man harms his foes thus: those things that they most dread he discovers, carefully investigates, then inflicts upon them.' Thucydides. E.g. Themistocles' use of spies to trick Xerxes into fighting the Battle of Salamis.

[210] 'If there was stiff competition round the centres of power, there was practically none in the area where I wanted to work – preparing the future.' Monnet (creator of the EU) assiduously targeted layers of advisors beneath the most senior people - those who do the crucial work during a 'beneficial crisis'. His small team repeatedly outmanoeuvred the UK Foreign Office.

[211] I have found great aversion to the introduction of checklists. It seems that people have an instinctive dislike of the idea that outsiders might reduce their expertise to some simple rules ('it's much more complicated than *that*,' they say.

patients and offer only problems for doctors (like their hands being burned by the disinfectants). Diarrhoea was in the anaesthesia category until the late 20th Century. Childbirth in the developing world remains in the anaesthesia category - we know what works but the ideas have not spread. *Technical solutions alone do not work; norms must change and this requires people talking to each other and using simple reliable tools like checklists.* (Gawande, 2013, see Endnote on Failure)

- New forms of training. For example, after the shock of the Yom Kippur War, Israel established the '*Talpiot*' programme which 'aims to provide the IDF and the defense establishment with exceptional practitioners of research and development who have a combined understanding in the fields of security, the military, science, and technology. Its participants are taught to be mission-oriented problem-solvers. Each year, 50 qualified individuals are selected to participate in the program out of a pool of over 7,000 candidates. Criteria for acceptance include excellence in physical science and mathematics as well as an outstanding demonstration of leadership and character. The program's training lasts three years, which count towards the soldiers' three mandatory years of service. The educational period combines rigorous academic study in physics, computer science, and mathematics alongside intensive military training... During the breaks in the academic calendar, cadets undergo advanced military training… In addition to the three years of training, Talpiot cadets are required to serve an additional six years as a professional soldier. Throughout this period, they are placed in assorted elite technological units throughout the defense establishment and serve in central roles in the fields of research and development' (IDF, 2012). The programme has also helped the Israeli hi-tech economy.[212] Tetlock has begun his *Good Judgement Project* to see whether people can be trained to be better at forecasting: 'We now have a significant amount of evidence on this, and the evidence is that people can learn to become better. It's a slow process. It requires a lot of hard work, but some of our forecasters have really risen to the challenge in a remarkable way and are generating forecasts that are far more accurate than I would have ever supposed possible from past research in this area.' This is part of IARPA's ACE programme to develop aggregated forecast systems and crowdsourced prediction software. IARPA also has the SHARP programme to find ways to improve problem-solving skills for high-performing adults.
- 'Artificial immune systems' to mimic the success of evolutionary principles: (i) generate lots of solutions with random variation (try lots of different things), and ii) differential replication of the most effective agents (reinforce success, weed out failure).
- Mass collaborations.[213]

---

[212] The old Technical Faculty of the KGB Higher School (rebaptised after 1991) ran similar courses; one of its alumni is Yevgeny Kaspersky, whose company first publicly warned of Stuxnet and Flame. It would be interesting to collect information on elite intelligence and special forces training programmes (why are some better than others at decisions under pressure and surviving disaster?). E.g. Post-9/11, US special forces (acknowledged and covert) have greatly altered including adding intelligence roles that were previously others' responsibility or regarded as illegal for DOD employees. How does what is regarded as 'core training' for such teams vary and how is it changing?

[213] Cf. DARPA's 2009 Network Challenge, a prize to the first team of volunteers able to locate 10 large red balloons hidden around the country - a classic intelligence problem not solvable quickly by conventional means. Cf. '*Reflecting on the DARPA Red Balloon Challenge*', Cebrian *et al*. This article explains some of the problems with crowdsourcing, including sabotage, that emerged with a later DARPA challenge (the 'Shredder Challenge'). In 2013 DARPA successfully experimented with online tools to allow geographically dispersed teams to crowdsource design challenges for financial prizes. Tools like Ushaidi provide a free platform for crowdsourced information sharing on any project, from disaster response to education. Kickstarter is a platform for crowdsourced start-up funding ('crowdfunding'). See section on Nielsen above.

- Simulations and 'agent-based models'. The [FutureICT project](#) is arguably the most ambitious project to simulate society, combining huge data inputs, mass collaboration, agent-based models, and simulations.[214] See Section 7.
- Visualisations.[215]
- Wargames and '[red teams](#)'.[216] (Brainstorming does not work.[217]) Simulations for training surgeons and others already exist. There are already various [programmes](#) to use online games for training military and intelligence personnel, such as 'NeuroTracker' used by NFL teams and special forces, and game avatars using real data (e.g. strength, speed, skill). The US Army is [experimenting with immersive virtual reality game training](#) and uses customized modules of the open-source computer game 'Unreal Tournament' for training. IARPA's 'SIRIUS programme' is developing games to train intelligence analysts: 'The goal ... is to create Serious Games to train participants and measure their proficiency in recognizing and mitigating cognitive biases... The research objective is to experimentally manipulate variables in Virtual Learning Environments (VLE) to determine whether and how such variables might enable player-participant recognition and persistent mitigation of cognitive biases. The Program will provide a basis for experimental repeatability and independent validation of effects, and identify critical elements of design for effective analytic training in VLEs. The cognitive biases of interest that will be examined include: (1) Confirmation Bias, (2) Fundamental Attribution Error, (3) Bias Blind Spot, (4) Anchoring Bias, (5) Representativeness Bias, and (6) Projection Bias.'[218] Another interesting possibility is to use virtual worlds or massively multiplayer online games (MMOGs) as a locus for real-world games, not least since they often have 'terrorist' problems themselves, or to overlay MMOGs on the real

---

[214] There is a combination of the Planetary Nervous System (to measure the physical and social state of the world), the Living Earth Simulator (to simulate what-ifs), and the Global Participatory Forum ('an open data and modelling platform that creates a new information ecosystem that allows you to create new businesses, to come up with large-scale cooperation much more easily, and to lower the barriers for social, political and economic participation'). 'Just imagine Google Earth or Google StreetView filled with people, and have something like a serious kind of Second Life. Then we could have not just one history; we can check out many possible futures by actually trying out different financial architectures, or different decision rules, or different intellectual property rights and see what happens.' (Hellbing) See Section 7 on 'agent-based model' simulations.

[215] E.g. 'Augmented reality' displays allow the projection of vein networks onto the skin for surgery, of CAD designs onto real buildings, or of information onto spectacles (and soon contact lenses). Google Glass has already been modified, against Google's wishes, to run facial recognition software and 'Winky', a camera operated by a blink of the eye. The 'Minority Report' special effects are [becoming real](#) and various people are working on [3D holograms](#). After overhearing Israeli air force officers, some neurosurgeons [applied military visualisation](#) software to transform a 2D picture into a 3D simulation.

[216] Warren Buffett has proposed institutionalising Red Teams to limit damage done by egomaniac CEOs pursuing flawed mergers and acquisitions: 'it appears to me that there is only one way to get a rational and balanced discussion. Directors should hire a second advisor to make the case against the proposed acquisition, with its fee contingent on the deal not going through'. Of course, wargames and Red Teams are only useful if listened to. E.g. Between the world wars the Germany Army examined British exercises with armoured divisions and asked themselves, 'how might this affect future war?' and insights helped develop Blitzkrieg. Japan's wargaming before Pearl Harbor assumed carriers would continue to be peripheral and in its planning for Midway, Rear Admiral Ugaki repeatedly overruled umpires whenever they made a decision that cast doubt on the ability of the Japanese navy to execute its plans. Classified Pentagon wargames 1963 - 1965 (the SIGMA series) predicted that the main elements of US strategy in Vietnam would fail. They were ignored. The report of the *Senate Select Committee on Intelligence* re the CIA, Iraq and WMD concluded: 'The presumption that Iraq had active WMD programs was so strong that formalized mechanisms established to challenge assumptions and 'group think', such as 'red teams' ... were not utilized.'

[217] One of the most widely used tools in politics (and advertising) is shown to *reduce* performance. A partner at BBDO wrote '*Your Creative Power*' in 1948 which introduced 'brainstorming' - a group churning out lots of ideas without negative feedback. Repeated experiments over decades have shown that the same number of people working alone will do better at solving problems.

[218] In March 2013, IARPA [announced](#) it is examining Alternative Reality games.

world. Suarez's *Daemon* is a brilliant fictional example of how this might be done.[219] DARPA is [exploring](#) how to give soldiers online training that can give them an understanding of how foreign locals might behave: 'how computers can automatically construct interactive, cultural training models from the combined experiences of warfighters… [We need] a computer system that can automatically parse and aggregate people's stories about a given topic and reconstruct variations of those experiences. The outputs are interactive training simulations similar to role-playing videogames or choose-your-own-adventure books… Riedl's training-generation system is called "Scheherazade" after the queen in *One Thousand and One Nights* who saved her life by telling stories. In response to a user-provided topic, Scheherazade uses crowdsourcing technology to capture stories about a topic from multiple individuals and combines the stories into an integrated situational model that allows the user to make decisions and explore outcomes.' (Cf. '*Story Generation with Crowdsourced Plot Graphs*', Riedl *et al* (2013).)

- Information markets.[220] E.g. Tetlock's *Good Judgement Project* (see above) has also examined the use of information markets: 'It's hard to argue with the overall performance of the two markets over the first several months of forecasting. The prediction markets have outperformed the simple average of predictions from forecasters in each of our survey-based experimental conditions… Only our "super-forecasters" (drawn from the top 2% of all participants in Season 1) have proven to be consistently more accurate than the prediction markets.'
- Network analysis. Cf. The work of [Barabasi](#) and companies such as [Palantir](#).
- Randomised trials.[221]
- Prizes (see Endnote.)
- Research (including about the sphinx of mass psychology[222]).

---

[219] One can imagine an interesting programme bringing together: a) historical case studies (e.g. the creation of the EEC/EU); b) online databases of reliable historical information (many sources use inconsistent information); c) realistic strategy games. Such a programme might not only train decision-makers much better but also make contributions to scholarship and make subjects fascinating to more children.

[220] E.g. The Iowa Electronic Market has predicted Presidential campaigns and other political conflicts better than pundits and polls (cf. Vernon Smith's [Nobel lecture](#)). Many companies use internal information markets, e.g. to harness information among staff to predict whether the firm will hit product launch timetables.

[221] James Lind did a trial for the Royal Navy in 1747 to test different diets' effects on scurvy and found that oranges and lemons worked. Randomised trials have since been used in countless ways. Cf. '*Supercrunchers*' and '*Adapt*'.

[222] Sometimes one may decide not to worry too much about how the public will react, for strategic, aesthetic or other reasons. It is important to know this is what you are doing. Generally, if you don't do market research you only have hunches about what normal people outside your world think and why. However, market research can kill creativity and original ideas, and can inhibit teams from breaking new ground. Market research was famously used by Irving Thalberg to produce hits, but it was also used to justify destroying Orson Welles' film *The Magnificent Ambersons*. After some focus groups in Pasadena, those saying things like 'It should be shelved as it is a crime to take people's hard-earned money for such artistic trash' and 'it's as bad if not worse than Citizen Kane' outweighed those who said, 'It seemed too deep for the average stupid person. I was disgusted with the way some people received this picture, which truly brings art to the picture industry.' RKO got in producer Bryan Foy who pronounced: 'Whole damn thing, too fuckin' long. Ya gotta take out forty minutes… Well, just throw all the footage up in the air and grab everything but forty minutes – it don't matter what the fuck you cut.' The year before, focus groups of *The Grapes of Wrath* had been awful but Darryl Zanuck said, 'Ship it - don't touch a foot.' There is no doubt that methods of empirical testing are improving but such judgements remain extremely difficult. In 1986, Martin *et al* published a twin study (*Proc. of National Academy of Sciences 83*, 1986) suggesting genes affect attitudes to political issues such as crime. It was almost totally ignored and was 'like a stone down a well' (Martin). In 2005, Hibbing *et al* published a similar study (*American Political Science Review 99*, 2005) also showing strong correlations between genes and political views. As with height, intelligence and other things, attention is now turning to GWAS.

• New and different kinds of institutions, like DARPA and Skunkworks.[223]

These tools should be applied by *new and different kinds of political teams*. Success will come to those who build new teams adding expertise from outside traditional political organisations[224] - for example, teams with the problem-solving abilities of physicists, understanding of cognitive science, advanced programming skills and so on - and who study and train in the art of making decisions under pressure.[225] In the 19th Century, the Prussian General Staff under the elder Moltke instituted new forms of training and planning, adapted faster than others to economic, technological and military trends, and, wielded by a politician of extremely unusual abilities, it shattered the established system. In the 20th Century, totalitarian movements adapted faster than traditional elites to developments in modern media and communications and, harnessed by leaders with great psychological insights about how to exploit them, shattered the established system. We must try to avoid similar developments in which dangerous individuals and entities emerge and adapt more effectively than our defences.[226] This means partly escaping a conundrum: we need institutions that operate like a successful immune system - that adapt faster than the enemy - yet bureaucracies, especially with democratic 'oversight', can rarely do this (and immune systems themselves can go haywire).[227]

We must also resist the temptation to think that extraordinary technical possibilities can solve fundamental ancient problems or remove the need for ancient lessons: as Colonel Boyd used to

---

[223] DARPA is the world's most successful government technology developer (e.g. stealth, the internet). It operates on different principles than normal state agencies: e.g. it is exempt from many regulations, it is explicitly high risk so fear of failure is limited, its program managers only stay a few years so there is a turnover of people re-examining assumptions and budgets, and it does not maintain its own labs and staff but outsources. IARPA is modeled on DARPA. Lockheed's famous Skunk Works was *separate* from the larger company with its own culture of excellence, speed, and agility. It produced the U2, Blackbird, and stealth bomber. The Santa Fe Institute, set up by Gell Mann, has pioneered transdisciplinary research in complex systems. JASON succeeded in bringing top scientists together with policy makers. The old East India Company pioneered appointment by competitive exam, created its own teaching college, and recruited leading intellectuals.

[224] Some, such as the Obama campaign, are now investing huge efforts in applying lessons from 'big data' to politics; e.g. using Randomised Control Trials to figure out exactly what best persuades people to vote or switch candidate.

[225] There are various dilemmas involved in recruiting people of high ability. This interview with a CIA psychologist explores some of the problems inherent in building teams with extremely focused people.

[226] The extreme centralisation of the UK media makes it vulnerable to cheap and effective 'information operations'. Much less than £1 million can, spent right, have dramatic and hard-to-detect effects.

[227] British PMs should appoint ministers from outside the Commons, put them in the Lords, and *give Lords 'rights of audience' to speak in the Commons* while simultaneously strengthening the Select Committee system (market research shows this is popular with the public). This would be a simple way to escape the problem of picking Secretaries of State expected to manage complex projects from the limited talent pool of the Commons. Arguably Britain should move to a directly elected PM and explore limits on the power of political parties including the end of taxpayer funding for parties (and term limits?). *The most important change in Whitehall is HR rules*: until it is possible to replace people quickly, in the same way that healthy armies fire bad generals fast, and *close whole departments and organisations*, major improvements are impossible. (E.g. Many in Whitehall said that reducing the DfE headcount substantially would make services worse but *performance has increased as headcount has decreased*, and I think that it would improve further if there were a further large decrease in the headcount towards a DfE of hundreds not thousands.) Embedding people with advanced modelling and statistical skills as part-time advisers to ministers would force Ministers and officials to confront the issue of acting on, against, or in the absence of, evidence. Allowing Ministers to hire specific project management people from outside Whitehall, and without Whitehall HR, could greatly improve quality of management. Massive changes are vital to a) our relationships with the EU and ECHR, b) judicial review, c) the role of lawyers and legal advice in Whitehall. There are hard trade-offs between the benefits and costs of transparency and scrutiny. Some parts of government need much more of both (including a stronger Freedom of Information Act), others much less. The US does not have EU procurement but it does have Congress. After 9/11, JSOC (the classified 'joint special forces command' responsible among other things for global strikes on Al Qaeda) established the *Special Capabilities Office*, a classified technology developer operating outside all laws and scrutiny. If it did not exist, one of those involved said, 'nothing would have got out to the field. Nothing. We exist because we have to exist' (Ambinder, 2013). JSOC reportedly often uses civilian networks, safes, and paper documents (easily burned) to avoid FOI requests.

shout, 'People, ideas, machines - in that order.'[228] There is a tendency to believe that education and technology can eradicate problems caused by human nature and Clausewitz's 'friction' and 'fog'. Problems can be temporarily limited but not eliminated; they are inherent in the nature of humans and of conflict (even unconscious bacteria use trickery). Formal education can only get future leaders and 'strategists' (in the proper sense of the word) so far. For some roles, intelligence and education alone are impotent - 'a revolution is not a dinner party or writing an essay' - while ambition without the training of the mind and character brings disaster.

*

---

[228] No technical developments can make up for ignoring ancient lessons: e.g. as of 2004, the State Department had <8 people classed as able to speak Arabic like a native. Any leader who thinks that drones mean spies in Afghanistan or Pakistan do not have to read Kipling's *Kim* any more will suffer many nasty surprises.

7. Political economy, philosophy, and avoiding catastrophes.

'What is government itself but the greatest of all reflections on human nature?' James Maddison.

'Thus revolution gave birth to every form of wickedness in Greece. The simplicity which is so large an element in a noble nature was laughed to scorn and disappeared... In general, the dishonest more easily gain credit for cleverness than the simple do for goodness; men take pride in one, but are ashamed of the other... At such a time, the life of the city was all in disorder, and human nature, which is always ready to transgress the laws, having now trampled them under foot, delighted to show that her passions were ungovernable, that she was stronger than justice, and the enemy of everything above her... When men are retaliating upon others, they are reckless of the future and do not hesitate to annul those common laws of humanity to which every individual trusts for his own hope of deliverance should he ever be overtaken by calamity; they forget that in their own hour of need they will look for them in vain... The cause of all these evils was the love of power, originating in avarice and ambition, and the party-spirit which is engendered by them when men are fairly embarked in a contest... For party associations are not based upon any established law nor do they seek the public good; they are formed in defiance of the laws and from self-interest...' Thucydides III.

'[Our] empire was not acquired by force... The development of our power was originally forced upon us by circumstances; fear was our first motive; afterwards honour, and then interest stepped in. And when we had incurred the hatred of most of our allies; when some of them had already revolted and been subjugated, and you were no longer the friends to us which you once had been, but suspicious and ill-disposed, how could we, without great risk, relax our hold? For the cities as fast as they fell away from us would have gone over to you. And no man is to be reproached who seizes every possible advantage when the danger is great.' Athenian envoys in Sparta.

'[O]f men we know that by a law of nature wherever they have the power they always rule. This law was not made by us and we are not the first who have acted upon it; we inherited it and expect to bequeath it to all time...' Athenian ambassadors to Melos, Thucydides V.

'If people do not believe that mathematics is simple, it is only because they do not realize how complicated life is.' von Neumann.

'Theoretically, planning may be good etc - but nobody has ever figured out the cause of government stupidity - and until they do and find the cure, all ideal plans will fall into quicksand'. Richard Feynman.

A growing fraction of the world has made a partial transition from (a) small, relatively simple, hierarchical, primitive, zero-sum hunter-gatherer tribes based on superstition (almost total ignorance of complex systems), shared aims, personal exchange and widespread violence, to (b) large, relatively complex, decentralised, technological nonzero-sum market-based cultures based on science (increasingly accurate predictions and control in some fields), diverse aims, impersonal exchange, trade, private property, and (roughly) equal protection under the law.[229]

How realistic is it to imagine a world in ~2050 in which a combination of science and market institutions enable nine-tenths of ~10 billion people to thrive in prosperous and peaceful nonzero-sum societies, fed by continued agricultural productivity growth (while reducing / reversing the impact of agriculture) and enjoying ubiquitous clean energy, clean water (extracted from the *oceans*), homes, healthcare, clean environments (fisheries and oceans recovering), and access to good education (*via* whatever the internet has become), as more of humans' collective efforts and

---

[229] Between hunter-gatherers and Manhattanites, there is a difference of ~400 times in income and ~30 million times in the range of products (~300 viz ~10 billion). Cf. Beinhocker, *Complexity Economics*.

intelligence shift from survival problems to some of the fundamental problems sketched here and tapping the effectively infinite resources of space?[230]

*

Why markets work

What we have learned about our world vindicates the *evolutionary* perspective of the pre-Socratics (Anaximander, Heraclitus), Thucydides, Hume, Smith,[231] Darwin, and Hayek[232] over the *anthropocentric* perspective of Plato, Aristotle,[233] Descartes,[234] Hobbes, Rousseau ('the general will'),[235] Bentham,[236] Mill (who introduced the concept of the 'natural monopoly'[237]) and Marx. Evolutionary biology, neuroscience, cognitive science, and behavioural genetics have undermined the basis for Descartes' *Ghost in the Machine*, Locke's *Blank Slate*, and Rousseau's *Noble Savage* and have established a scientific basis for exploring a universal human nature. Economic theory, practice, and experiment have undermined the basis for Cartesian central planning: decentralised coordination *via* market prices is generally a better method for dealing with vast numbers of

---

[230] Consider the divergent paths of places like the Congo and Singapore. They had similar GDP per capita 50 years ago. Congo has got worse, Singapore has overtaken Western Europe.

[231] Re Smith: 'For the first time it was shown that an evident order which was not the product of a designing human intelligence need not therefore be ascribed to the design of a higher supernatural intelligence, but that there was a third possibility - the emergence of order as the result of adaptive evolution' (Hayek). Darwin was influenced by Smith's invisible hand.

[232] 'The curious task of economics is to demonstrate to men how little they really know about what they imagine they can design. To the naive mind that can conceive of order only as the product of deliberate arrangement, it may seem absurd that in complex conditions order, and adaptation to the unknown, can be achieved more effectively by decentralizing decisions and that a division of authority will actually extend the possibility of overall order. Yet that decentralization actually leads to more information being taken into account.' 'To understand our civilization one must appreciate that the extended order (of cooperation) resulted not from human design or intention but spontaneously: it arose from unintentionally conforming to certain traditional and largely *moral* practices many of which men tend to dislike, whose significance they usually fail to understand, whose validity they cannot prove, and which have none the less fairly rapidly spread by evolutionary selection - the comparative increase in population and wealth - of those groups that happened to follow them…' Hayek.

[233] Aristotle's teleological attitude to nature: 'In some birds the legs are very long, the cause of this being that they inhabit marshes. I say the cause, because nature makes the organs for the function, and not the function for the organs' (*On the Parts of Animals, IV, 12*). '[E]xperience shows that a very populous city can rarely, if ever, be well governed… For law is order … but a very great multitude cannot be orderly: to introduce order into the unlimited is the work of a divine power – of such a power as holds together the universe… [A] state when composed of too few is not, as a state ought to be, self-sufficient; when of too many, though self-sufficient in all mere necessaries … it is not a state, being almost incapable of constitutional government. For who can be the general of such a vast multitude, or who the herald, unless he have the voice of a Stentor?' (*Politics*).

[234] Descartes' epistemological scepticism (*De omnibus dubitandum*, *Everything is to be doubted)* made him the father of science and 'thus the grandfather of the Revolution' (Nietzsche). However, Descartes thought that the best things, whether city design or law-giving, are the product of a single mind consciously executing a plan - not messy evolution. '[T]he peoples who were formerly half savages, and who became civilised gradually, making their laws only in so far as the harm done by crimes and quarrels forced them to do so, could not be so well organised as those who, from the moment at which they came together in association, observed the basic laws of some wise legislator.'

[235] Locke *contra* Rousseau: 'The freedom then of man, and liberty of acting according to his own will, is grounded on his having reason, which is able to instruct him in that law he is to govern himself by, and make him know how far he is left to the freedom of his own will. To turn him loose to an unrestrained liberty, before he has reason to guide him, is not the allowing him the privilege of his nature to be free; but to thrust him out amongst brutes.'

[236] Bentham wished to remake all institutions on 'rational' principles. Von Neumann said of Bentham's famous idea that we should pursue 'the greatest possible good for the greatest possible number': it is 'taken literally, self-contradictory... A guiding principle cannot be formulated by the requirement of maximising two (or more) functions at once.'

[237] Mill thought it wasteful to have two carriers operating on the same route, could not imagine two cities connected by parallel railroad tracks, and is 'the intellectual father' of various monopolies around the world (Smith, 2002).

possibilities than Cartesian or Soviet planning, though obviously markets have problems particularly with monetary policy and financial regulation.

Vernon Smith won the economics Nobel Prize for his pioneering work on testing economy theory with laboratory experiments. Researchers have compared i) standard 'rational expectations' models and ii) thousands of experiments of a) *personal exchange* or b) *impersonal exchange* (i.e. 'market experiments where cooperation can occur through the coordination function of prices produced by, but simultaneously resulting from, interaction with individual choice behavior'). The conclusions: first, humans perform *better* than predicted in many personal exchange experiments such as two-person iterated Prisoners' Dilemma games (even if anonymous) because they cooperate more than standard models predict.[238] Second, overall, in *impersonal exchange* experiments with many different types of institution, self-interested non-cooperative subjects trading in markets (which are often too complicated for 'formal game theoretic analysis to articulate predictive models') generate efficient outcomes and overall gains close to the predictions of the standard model; people's 'autonomic mental algorithms coordinate behavior through the rules of the institution' (i.e. 'social algorithms') and promote social ends that are not part of their intention, in accord with Smith's invisible hand.[239]

The results of experiments using the 'continuous bid/ask double auction' (CDA), in which different types of participants consistently converge *via* trading on a solution of nonlinear equations from standard theory (that participants could not solve mathematically, or usually even read), are particularly striking:

> 'What we learn from such experiments is that any group of people can walk into a room, be incentivized with a well-defined private economic environment, have the rules of the oral double auction explained to them for the first time, and *they can make a market that usually converges to a competitive equilibrium, and is 100 percent efficient - they maximize the gains from exchange* - within two or three repetitions of a trading period. Yet knowledge is dispersed,

---

[238] Trivers' theory of evolved 'reciprocal altruism' is thought to lie behind the 'tit-for-tat' strategy we observe in humans and computers playing Iterated Prisoners' Dilemma games. 'Genetic algorithms' playing each other in computer-based tournaments (where game-playing strategies evolve according to evolutionary principles depending on their success) also evolve variations of tit-for-tat, which has been widely observed in biology (e.g. shoals of fish approaching predators). Cf. Axelrod's classic '*The Evolution of Cooperation*'. *In July 2012, dramatic new results for the Prisoners' Dilemma were announced after a game theorist emailed Freeman Dyson: see* here, *and papers* here *and* here. Recent work has found 'generous zero-determinant strategies that tend to dominate in evolving populations … [and] are not only evolutionarily stable, but they also can invade and replace any other ZD strategy [and] are robust to invasion by any non-ZD strategy… [W]hen it comes to evolutionary success, it is generosity, not extortion, that rules' (*From extortion to generosity, the evolution of zero-determinant strategies in the prisoner's dilemma*, Stewart & Plotkin (2013)). Cf. *How Natural Selection Can Create Both Self- and Other-Regarding Preferences, and Networked Minds, Grund et al (Nature, 2013)*.

[239] There are many examples of humans *not* conforming to the predictions of the standard models but this is not surprising since the research programme of behavioural economists has been a 'deliberate search in the tails of distributions', and instead of just labeling people 'irrational' it is better to probe into *why* people behave as they do and what sort of problem they think they trying to solve (V. Smith). 'The claims that people are not good Bayesian intuitive statisticians has been considered one more source of human cognitive errors… But such predictive errors can just as strongly speak to weaknesses in the theory as to flaws in human judgement under uncertainty. An alternative to the Bayesian hypothesis is that human decision makers are adapted and attuned to decision-making in uncertain environments that are not characterised by the ball-in-urns sampling paradigm… *[T]he human brain might be trying to solve a different and more relevant problem* than the one posed by the Bayesian model… I think a good working hypothesis is the following: We owe our existence to ancestors who developed an unusual capacity to adapt to surprises, rather than to process sample information optimally from a stable historical environment in which all states were specifiable and presumed known [p.154ff]…The Bayesian model postulates that a person can a priori list all possible states of nature… Life experiences are analagous to sampling an urn that at a given time is perhaps thought to contain only red and black balls, but a new event experience yields a new outcome - an orange ball… How can it be fit, as in Bayesian modelling, for any well-adapted organism to assign a zero posterior probability to any event surprise merely because there was no prior experiential basis for a positive assignment?' (Smith p.321)

with no participant informed of market supply and demand, or even understanding what that means.'[240] (Emphasis added)

Further, these results also show that results consistent with standard theoretical predictions also occur under *weaker* conditions than the conventional theory requires: there need *not* be a large number of agents, they do *not* need 'perfect information' or be economically sophisticated. These results are, therefore, 'a giant victory for the economic theory of markets', but they are also a demonstration that 'the theory is incomplete' and 'we do not understand why markets work as they do.'

> 'What experimentalists have unintentionally brought to the table is a methodology for objectively testing the Scottish-Hayekian hypotheses under scientific controls. This answers the question Milton Friedman is said to have raised concerning the validity of Hayek's theory/ reasoning: 'how would you know?'… The pattern of results greatly modifies the prevailing, and I believe misguided, rational SSSM ['standard social science model'], and richly modernizes the unadulterated message of the Scottish philosophers.' (V. Smith, op. cit)

The position of Hume, Adam Smith, and Hayek is 'the antithesis of the anthropocentric belief that if an observed social mechanism is functional, somebody in the unrecorded past must have used reason consciously to create it to serve its perceived intended purposes.'

> 'Since our theories and thought processes about social systems involve the conscious and deliberate use of reason, it is necessary to constantly remind ourselves that human activity is diffused and dominated by unconscious, autonomic, neuropsychological systems that enable people to function effectively without always calling upon the brain's scarcest resource – attentional and reasoning circuitry. This is an important economizing property of how the brain works. If it were otherwise, no one could get through the day under the burden of the selfconscious monitoring and planning of every trivial action in detail… We fail utterly to possess natural mechanisms for reminding ourselves of the brain's offline activities and accomplishments…'[241]

> 'Just as most of what the brain does is inaccessible to the mind, so also is there a widespread failure of people to understand markets as self-organizing systems, coordinated by prices for cooperative achievement of gains from exchange, without anyone being in charge. The workings of the economy are as inaccessible to the awareness of its agents, business persons included, as is an agent's awareness of his own brain functioning. The workings of the economy are not the product, nor can they be the product, of conscious reason, which must recognize its own limitations and face, in the words of Hayek, "the implications of the astonishing fact, revealed by economics and biology, that order generated without design can far outstrip plans men consciously contrive".'

The combination of the findings concerning i) the evolved propensity of people to cooperate in *personal exchange* and ii) the success of *impersonal exchange* (markets) in maximising gains through trade helps explain one of the central problems of intellectual and political discussion of markets. *Experience* of personal exchange, where deliberate planning and cooperating works, makes people

---

[240] CDA is more efficient than 'posted offer' (i.e. retail pricing). The high cost of training staff to use CDA has prevented its widespread use but the internet is changing this. Cf. V Smith's [Nobel lecture](). The *Iowa Electronic Market* (IEM), a futures market for political bets, is an example of how the CDA demonstrates great results in the real world.

[241] Language and almost everything to do with our developmental socialisation is learned without explicit instruction. Amnesiacs can learn new skills and perform them without memory of having learned them.

think that they can successfully intervene in markets that actually work because they coordinate, *via* prices, vastly more complicated behaviour than individuals can analyse or plan effectively.

> 'In impersonal exchange through markets, the gains from exchange are not part of their experience. As noted by Adam Smith, 'This division of labour . . . is not originally the effect of any human wisdom, which foresees and intends that general opulence to which it gives occasion.' Impersonal exchange through markets tends to be perceived as a zero sum game, which perception in no way diminishes the capacity for markets to do the work articulated by Adam Smith... Interventionist programs ... result from people inappropriately applying their intuition and experience from personal social exchange to markets, and concluding that it should be possible to intervene and make things better. People use their intuition, not their reason ... in thinking about markets, and they get it wrong.'

Hayek described a market (based on legal protection by more or less independent courts of individuals' property 'equally under law') as an evolutionary *information-discovery process* in which price-information signals allow adaptation to an ever-changing complex environment; many ideas are tried, successes are reinforced and failures are weeded out. Enterprises bring together new information about particular circumstances (that is not available to distant, busy politicians or officials) and the specialisation of knowledge (usually known as 'the division of labour') for long-term (by political standards) activity. Markets are far from perfect: for example, [this memoir](#) provides a fascinating glimpse into how perhaps the best advert for markets on earth, Silicon Valley, has relied on government investment in basic science the consequences of which often (such as with the internet) eluded the most successful entrepreneurs for years. Individual entrepreneurs and firms do not necessarily learn - they can and do persist with doomed ideas - but the damage they do is contained, they go bust, and resources are diverted to new ideas, while bureaucracies can keep ignoring reality and keep ploughing more and more resources into disastrous enterprises. Markets allow a fragile truce between Apollo and Dionysus: reason has space to flourish, violent instincts are kept in check, but people have considerable freedom including to further the interests of relatives. Markets are non-zero-sum so they provide an alternative to violence (in-group or out-group) over scarce resources.[242]

Although some modern scholars have equated Adam Smith and other contemporaries with the idea that man is rationally self-interested in a purely economic sense, Smith himself had no such belief and was deeply concerned with questions of morality and motives in his *Theory of Moral Sentiments* which is now little read. The essential point of Smith, Mandeville *et al* was that doing good for others does not *require* deliberate action with the aim of doing good for others.[243] This

---

[242] An example of the evolutionary discovery-process of markets is the deregulation of airlines that led to a major reorganisation of the network of flight routes into what is now known as 'the hub and spoke' system. This reorganisation was predicted by nobody and was not part of the rationale for deregulation; *it evolved*. Smith argues that *it could not have been predicted* because the preferences on which the hub-and-spoke system was based had to be discovered through market experimentation. 'Unknown to both managers and customers was the subsequently revealed decision preferences of customers who unknowingly favored frequency of daily departure and arrival times – a preference that had to be discovered through market experimentation.' Often, examples of supposed 'market failure' are actually examples of the problems caused by bad regulation. E.g. the oft-cited California energy crisis in the 1990s. Lighthouses are famously used as an example of something that 'must' be a public good. 'It turned out that early lighthouses were private enterprise, not government, solutions to a public good problem, and the alleged inevitability of free riding was solved by the owner who contracted with port authorities to collect lighthouse fees when ships arrived portside to load or unload cargo.' V. Smith, 2002.

[243] 'He generally, indeed, neither intends to promote the public interest, nor knows how much he is promoting it… [B]y directing that industry in such a manner as its produce may be of the greatest value, he intends only his own gain, and he is in this, as in many other cases, led by an invisible hand to promote an end which was no part of his intention. Nor is it always the worse for the society that it was no part of it. By pursuing his own interest he frequently promotes that of the society more effectually than when he really intends to promote it.' Adam Smith.

has been mis-interpreted by many modern thinkers as meaning that markets *require, justify and promote selfishness*. 'People are not required to be selfish; rather the point of the Scottish philosophers was that people did not have to be good to produce good. Markets economize on information, understanding, rationality, numbers of agents, and virtue' (V. Smith). Understanding how humans really make decisions requires developments outside traditional economics, particularly theoretical and experimental information about the mind.[244]

<p align="center">*</p>

Science, technology and markets: Europe (open, competitive, innovative) versus China *et al* (closed, hierarchical, conservative)[245]

13th Century China was more advanced than Europe. They had already invented paper,[246] woodblock printing, gunpowder and paper currency. They had hundreds of thousands of educated people who competed, *via* sophisticated multiday exams, for administrative posts. Their two biggest cities had over a million people each, dwarfing Paris. By the 19th Century when China clashed with European powers, it was routed by numerically inferior but technologically and organisationally superior forces.

Why? Markets and science...

In England, the Great Famine and Black Death of the 14th Century and peasant revolts of 1381 encouraged a transition from serfs to tenant farmers with property rights. Growing agricultural productivity allowed greater division of labour and more manufacturing and innovation. From this period, England in particular developed institutions that enabled markets to thrive and these markets exploited growing scientific and technological knowledge. The combination of the Renaissance, the invention of the printing press in the 15th Century,[247] the birth of the scientific method from the early 17th Century, and breakthroughs from Galileo, Kepler, Descartes (algebraic geometry) *et al* provided foundations for Newton (calculus and mechanics) and the birth of the Industrial Revolution.

Meanwhile, China stagnated. The strong central control of the Chinese imperial bureaucracy contrasted sharply with the competition between European states and institutions and thus

---

[244] 'I importune students to read narrowly within economics, but widely in science. Within economics there is essentially only one model to be adapted to every application: optimization subject to constraints due to resource limitations, institutional rules and /or the behavior of others, as in Cournot-Nash equilibria. The economic literature is not the best place to find new inspiration beyond these traditional technical methods of modelling.' V. Smith, op. cit. For Hume, *contra* Descartes, 'the rules of morality … are not conclusions of reason'. However, we can use reason to explore the emergent order in human cultures: 'to discover the possible intelligence embodied in the rules, norms and institutions of our cultural and biological heritage that are created from human interactions but not by deliberate human design. People follow rules without being able to articulate them, but they can be discovered.' Experiments are now designed to test robot and human behaviour in economic situations with different institutional rules and humans can be connected to sensors such as fMRI scans (e.g. suppressing brain activity in some regions affects behaviour in the Ultimatum Game).

[245] Much of this section comes from '*The Infinite Resource*' by Naam (2013).

[246] Paper, invented by the Chinese ~7th Century, came to Europe via the battle of Talas in the 8th Century when an Arab army beat a Chinese army in central Asia and Chinese prisoners taught them to make paper. By the 1200s it reached Europe where it combined with other ideas, including Chinese woodblock printing, to give Gutenberg his idea of making the individual symbol the basic unit of printing in the 15th Century.

[247] The invention of the printing press in the 15th Century saw the price of a book fall rapidly by about 300 times from the equivalent of about $20,000 pre-Guttenberg to about $70, increasing the number of books by about 30 times in the first century after invention. This not only helped the spread of Shakespeare, maths, and science but also religious persecution (the Spanish Inquisition, the Thirty Years War etc). (Silver, 2012)

prevented the scale of market and intellectual innovation that rippled across Europe. While the Chinese bureaucratic exam system spread education,[248] it also (reinforced by state control of the press) imposed a rigid and static intellectual outlook (including Confucianism) while Europe benefited from dynamic intellectual revolutions.[249] Also, the Chinese imperial system had a strong system of state monopolies and successful new commercial ventures were often nationalised. Under the Ming dynasty, there was a general attitude that China needed nothing from the outside world which led to the infamous policy of *haijin* - the 'sea-ban'. Columbus traipsed around Europe for years pitching his ideas for exploratory voyages before eventually succeeding with Isabella of Spain in 1492. In China, 90 years before Columbus' voyage in the 15th Century, the Chinese emperor sent out a huge fleet of hundreds of vessels that reached east Africa. However, in 1424 the fleet was burned because of *haijin*. Admiral Zheng could not, as Columbus could, go to another source of money and support (or rather, if he had wanted to, it would have meant a death sentence so it would have been a much higher risk than for Columbus).

Similarly, in 1630s Japan, the Tokugawa shogun Iemitsu created the policy of *sakoku*, prohibiting, on pain of death, Japanese to leave or foreigners to enter; western books were burned, and Japan remained largely isolated until 1854. Although Islam had helped preserve classical knowledge from the ancient world, made their own innovations, and transported ideas from China and India (such as our modern number system including zero), the Ottoman Empire also suppressed rather than embraced the printing press and banned it from 1483 (with partial exceptions for foreigners).[250]

<div align="center">*</div>

The good news

More, richer, safer, educated people...

The spread of science, technology and markets has brought vast global improvements. Over the past two hundred years, the global population has risen fast, productivity (including agricultural productivity) has risen even faster, and the amount of energy used per person has risen. These billions are healthier, richer, and less likely to die from disease or violence.

The global population rose from ~1b in 1800[251] to ~2b in 1920s, ~4b in the 1970s, ~7 billion now, and perhaps ~8 billion by 2030. The average number of children per woman has halved over the last fifty years, from 4.9 children per woman to 2.5 in 2011 (the population is stable at 2.1). It is reasonable to predict that the population will stabilise at 9-10 billion ($\sim 10^{10}$) in the next few decades, with most of these new billions living in cities.

Life expectancy has grown from ~47 in 1950 to ~68 (~80 in the West) and the gap between developed and developing nations has shrunk from ~25 to ~13 years. In 1850, ~⅓ of children died by their fifth birthday; by1950, it fell to ~1/7 and by 2010 to ~1/20.

---
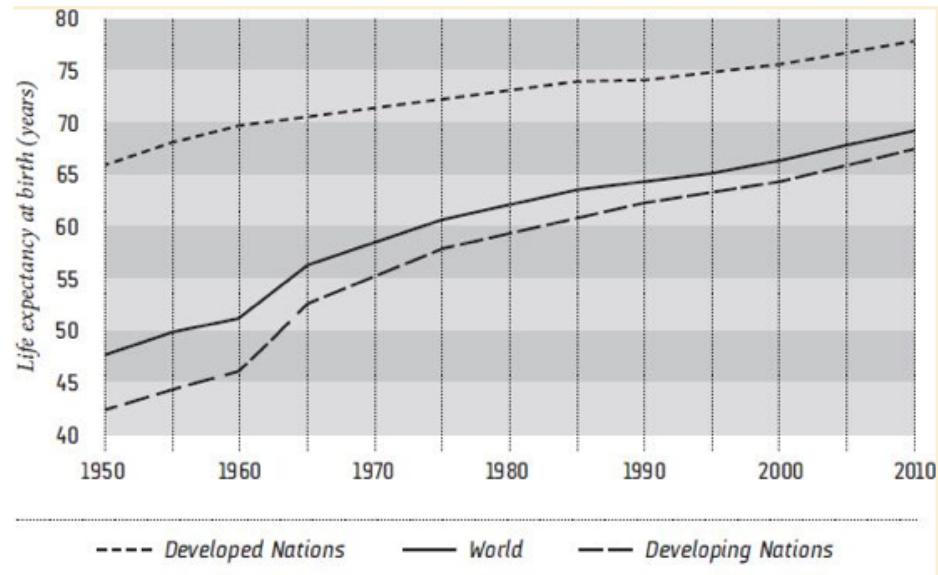
[248] Cf. '*Chinese Influence on The Western Examination System*', Teng (1943) for the history and influence of China's system.

[249] When Yingxing, a great scholar, produced his encyclopedia '*The Exploitation of the Works of Nature*', most copies were burned because it provided explanations of things the state wanted to keep control of.

[250] The effects of Lysenko's infamous anti-Darwinism on Stalinist Russia's biologists show the interaction between a society's openness/closedness, scientific development, and economy.
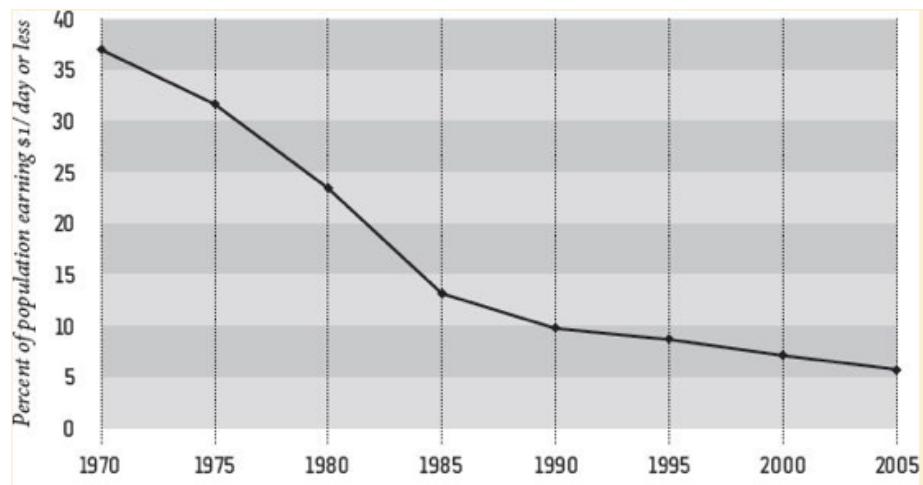
[251] ~350 million after the Black Death in the 14th Century.

Graph: Rising life expectancy

In 1970, over a third of the developing world lived on less than $1 per day; today, only ~5% do and the number is falling. In 1970, about half of the planet lived on less than $1,000 per year; today, less than a fifth do.

Graph: Falling poverty

Money goes further as productivity improves. Developments in agriculture over 10,000 years have reduced the amount of land needed to feed a person by a factor of ~1,000 to 10,000: farms in the developed world can now feed nearly 3,000 people per square mile. The number of people that the planet can feed has therefore grown dramatically and there is no reason why it cannot continue to grow. The price of food fell by a factor of 3 over the 20ᵗʰ Century. The total input of energy and water into agriculture has grown *less* than the amount of food produced: according to the USDA, *the amount of energy used per calorie has halved over the past 50 years*. This is due to the 'Green Revolution', pioneered by Norman Borlaug (funded by the Rockefeller Foundation and the Mexican government), which substantially increased agricultural yields. Water productivity has also risen significantly.
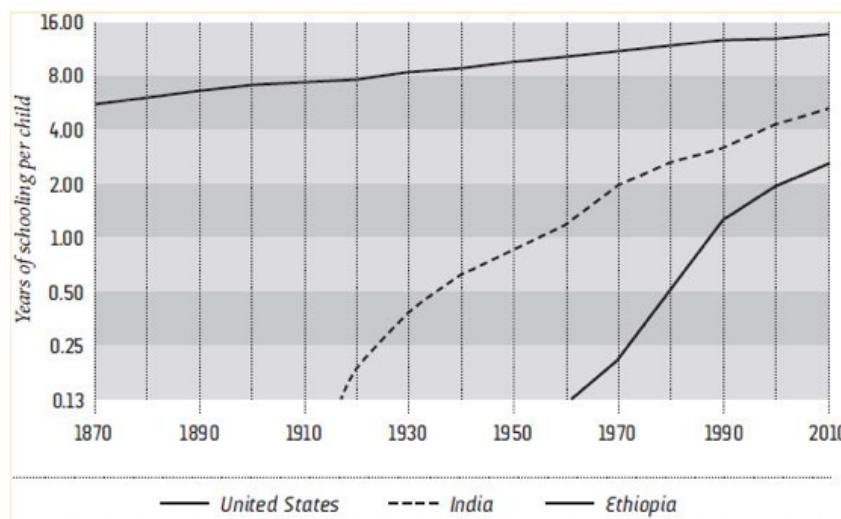
Overall, 1900-2000 the prices of metals, adjusted for inflation, fell by a factor of 4.5. In America, where prices of raw materials have fallen relative to inflation while incomes have risen relative to inflation, the gains are higher.

> 'The result is that, relative to average salaries in the United States, the price of food dropped by a factor of 25 and the price of metals by a factor of 35 across the twentieth century. The prices of energy, clothing, furniture, appliances, and electronics have all dropped when compared to income in the last decade, and in every decade since World War II.' (Naam)

Although Americans have spent about the same on housing over 60 years, the average size has more than doubled while the number of people per house has dropped by a quarter.

Education has spread across the world and, crucially, among girls.

Graph: growing and spreading education

In 1997, a fifth of the developed world had a mobile phone; now four-fifths of the developing world have one. Between 1990-5, internet access spread from a tiny number of scientists to a few million early adopters; now, about a third of the world has access, mostly mobile. We have done all these things while also *improving energy productivity* (cf. Section 2).

These improvements were accelerated by the transition, between 1977 and the early 1990s after the end of the Soviet Union, of China, the former Soviet Unions, India and much of the 'developing' world to market-based societies. Between 1977 and 1991, Deng changed China's path from Maoism to markets, the Soviet Union collapsed, and India opened to foreign companies. More than a third of the world shifted from state controlled backwardness to market-based societies. After 60 years, the market-based South Korea is about 20 times richer per capita than the mafia-state of North Korea with its endemic famines. South Korea was similar to Egypt in 1960 and is now five times richer per capita. Fifty years ago, the Congo and Singapore had similar GDP per capita; now Singapore is richer than Britain. Commercial cars, flight, penicillin, radio, TV, plastics, the computer, internet, and mobile phones all were developed in market-based America and Europe but the opening of Asia means the end of the West's dominance of science and technology.

Shortage and innovation...

There is a long history of shortage leading to price rises, exploitation, and innovation. Whale oil used for lighting by the rich was replaced first by the innovation of kerosene then by Edison's
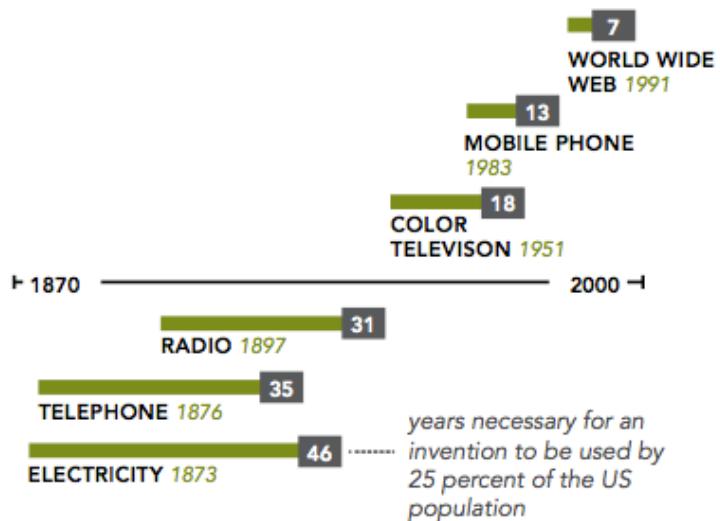
lightbulb. Humans used manure as a fertilizer for centuries without understanding anything about the chemistry.[252] In 1905, Fritz Haber showed how to produce ammonia direct from the atmosphere. The 'Haber-Bosch process' produces almost all synthetic fertiliser, uses about 1 percent of the world's energy, and roughly doubles the amount of food we grow. During World War II, the Allies produced synthetic rubber and the Germans produced synthetic gasoline (*via* the Fischer-Tropsch process).

It takes between 1,000 and 10,000 times less land to grow food than in antiquity; it takes 100 times less labour than a century ago; it takes 2-3 times less energy and water to grow grain than fifty years ago; over 150 years, the amount of energy needed to produce a unit of light has reduced by a factor of nearly 500 while the price has declined; we use 40-50 times less energy to turn a ton of iron into a ton of steel than in the 19[th] century; the amount of energy used to heat a US home has halved since 1978.

There are lots of things we are *not* short of: many raw materials such as iron and titanium are so abundant in the earth's crust that we have *millions of years* supply at current extraction rates. Already many materials are effectively recycled: over half the copper and paper and ~40% of steel used in the US comes from recycling. The use of recycled materials can reduce the amount of energy needed.

Innovations are also spreading faster.

Chart: Innovations spreading (Source: National Intelligence Council, 2013)...
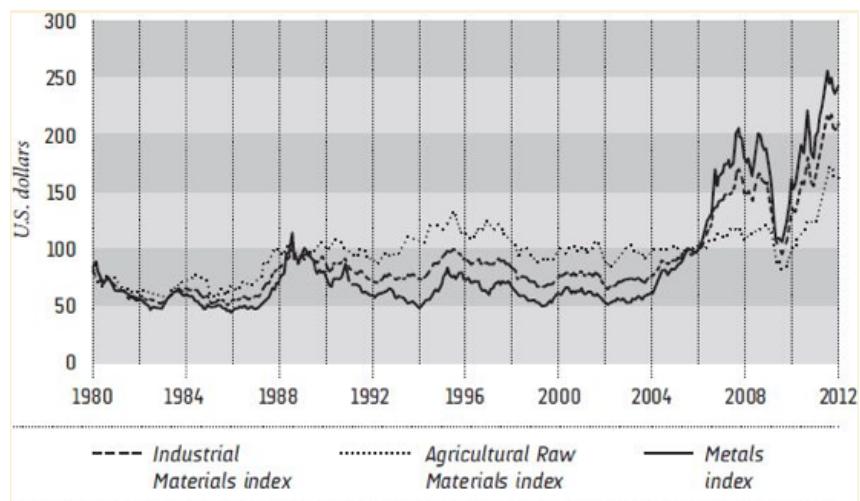
Some bad news...

Commodity and food prices have risen sharply recently, fish stocks have fallen, there are serious shortages of freshwater as the great aquifers shrink fast,[253] and deforestation triggers various negative feedbacks.

---

[252] In the 19[th] Century, there were various innovations such as Peru's guano then saltpeter (the discovery of both provoked wars).

[253] Less than 0.5% of the world's water is liquid freshwater and another 2% is in the polar ice caps and glaciers. If we can use solar energy to desalinate ocean water, then we can use the other 97%. The amount of energy needed to desalinate water has fallen by a factor of 9 since 1970. Already 10% of Singapore's water is desalinated.

Graph: Rising commodity prices

Industrial Materials index ---- | Agricultural Raw Materials index .... | Metals index ——

I will not go into any details of the 'global warming' debate. There has undoubtedly been hype and dishonesty from many people exaggerating the precision of our knowledge and distorting science for various purposes (anti-market political motives, funding and so on).[254] Obviously, detailed predictions of the same kind as particle physics are impossible. Many policies implemented, and spending decisions taken, 'because of global warming' (or for cynical PR reasons) are stupid, counterproductive, and a waste of money.[255]

However, this does not mean that the overall scientific consensus can or should be ignored as many argue. The basic physics of 'greenhouse gases' is not controversial among physicists. CO_2 molecules allow visible light from the sun to pass through them but they absorb infrared light. The planet glows infrared. CO_2, water vapour and other gases trap some of this infrared heat in our atmosphere. This effect is vital for life but it is also a risk. There is a widespread consensus among scientists that the planet is warming, that human activity which increases greenhouse gases is the main reason, and that a 'business as usual' approach would mean rises in CO_2 levels that would bring temperature rises of a scale that may have significant effects - some good but mainly damaging. How bad those effects may be and the probability of very bad case scenarios are impossible to answer with precision and model predictions regularly change (cf. the interesting chapter on climate modelling in Nate Silver's '*The Signal and The Noise*'). What we do about it is very controversial and *is not something that science alone can answer* but given the scale of possible problems the issue obviously cannot be ignored. Unfortunately, many on the Left want to use the issue to attack markets, many on the Right indulge an anti-scientific attitude (the fact the BBC exaggerates / distorts does not mean the issue can be ignored), many on all sides of politics use the issue to signal 'values', and some scientists have unwisely sacrificed long-term credibility for short-term impact.

Professor Robert Muller (Berkeley), an expert on climate change and main author of one of the recent studies of global temperatures, concludes:

---

[254] E.g. The main author of the IPCC 2007 Report made a terrible decision, on PR grounds, to leave in a prediction about the effects on India and China of the Himalayan glacier melting. Such things undermine trust in scientific advice.

[255] E.g 'The EU's mandate for biofuels has been as bad [as America's subsidies]. It's encouraged the use of palm oil, which also takes a tremendous amount of land to grow and other fossil fuels to refine into biodiesel. The result has been the deforestation of 6 million acres of rainforest in Indonesia, and the emission of more CO_2 than would have been emitted by burning diesel from traditional oil wells.' Naam (2013).
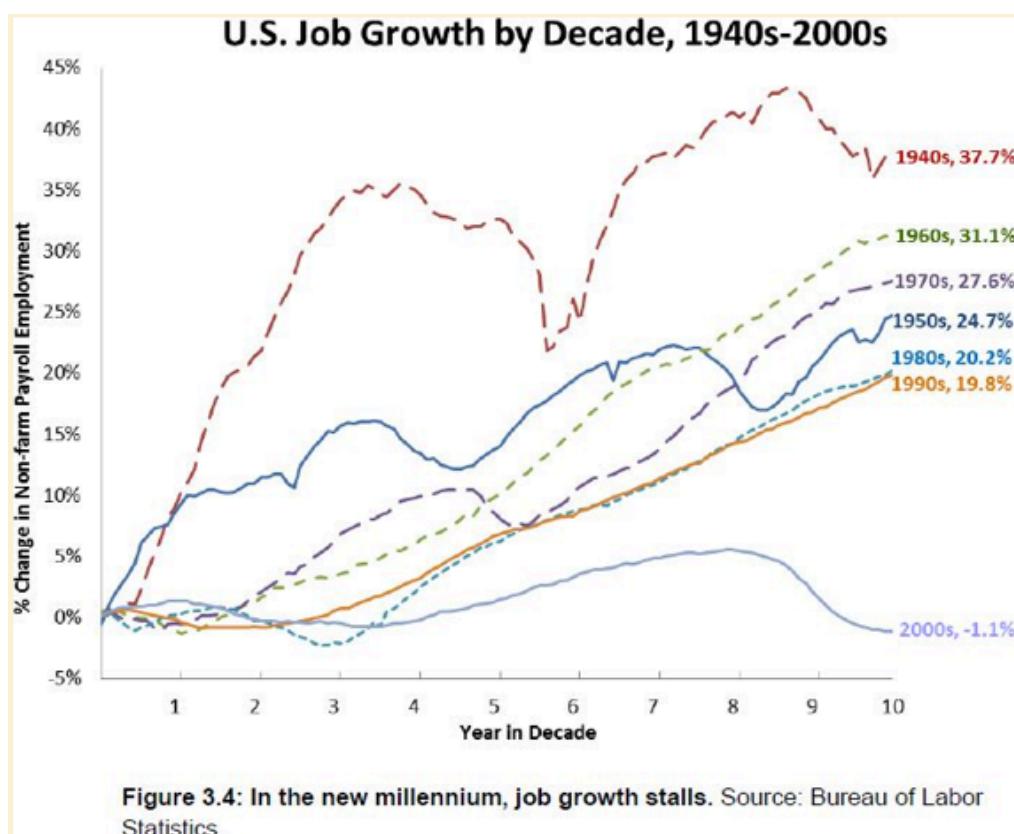
'Virtually none of the publicly proposed solutions to the danger of increased carbon dioxide, if implemented, have any realistic chance of working. "Setting an example" fails if the example is one that the developing world can't afford to follow. Perhaps the only workable solution is to encourage coal-to-shale-gas conversion in the developing world by vigorous sharing of US know-how.'

Is the economy stagnating…?

In America, there is concern about the lack of growth of employment and median income.
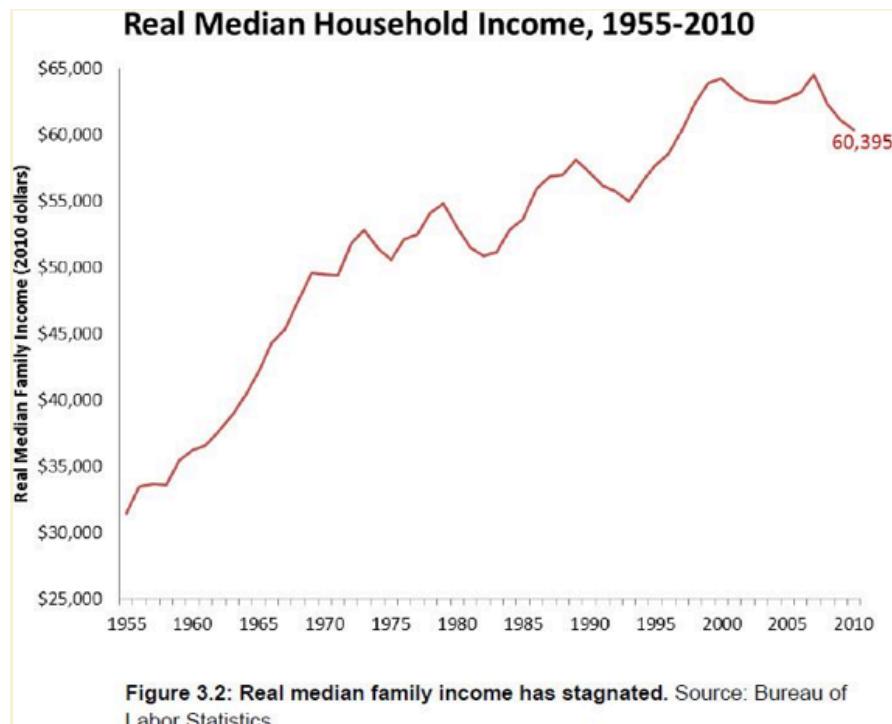
Less job growth… The workforce participation rate in the US has fallen to about two-thirds - a level not seen since 1983 when women had not yet entered the labor force in large numbers. Despite the population growing by 30 million over the decade 2000-2010, that decade was the first since the Great Depression to produce no net job growth.

Graph: Stagnant US job growth (Brynjolfsson)

Figure 3.4: In the new millennium, job growth stalls. Source: Bureau of Labor Statistics.

Stagnant median incomes… 'The growth of median income slowed down significantly at least 30 years ago, and actually declined during the first decade of this century; the average family in America earned less in 2009 than it did in 1999' (Brynjolfsson).
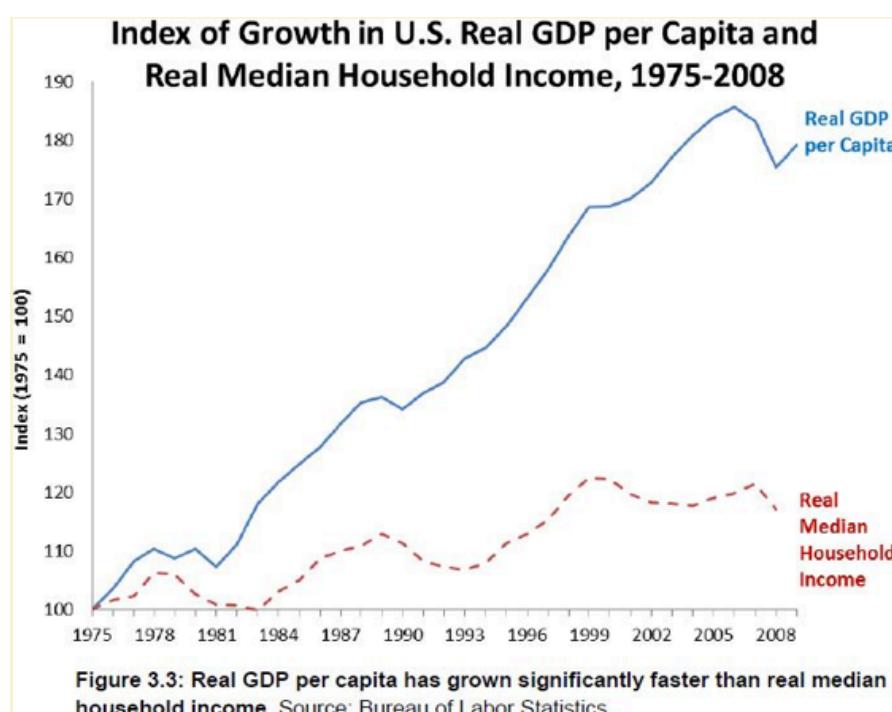
Graph: Real median income has stagnated (Brynjolfsson)

Real Median Household Income, 1955-2010

Figure 3.2: Real median family income has stagnated. Source: Bureau of Labor Statistics.

However, GDP per capita continued to rise.

Graph: GDP per capita compared with real median household income (Brynjolfsson)

Index of Growth in U.S. Real GDP per Capita and Real Median Household Income, 1975-2008

Figure 3.3: Real GDP per capita has grown significantly faster than real median household income. Source: Bureau of Labor Statistics.

There are three main arguments for current economic woes: *cyclicality, stagnation, and the 'end of work.'*

1. Cyclicality. Krugman argues that unemployment is the result of 'inadequate demand - full stop.'

2. Stagnation. Cowen and others argue that while the 2008 crisis obviously made things worse, there were already long-term structural changes afoot that will end of the levels of economic growth we have seen since 1750: we have been living off 'low-hanging fruit' for 300 years and are now struggling to be more productive, as declining productivity figures show.[256] Further, as markets have spread around the world competition has exerted a downward pressure on the growth of US employment and earnings. In 2012, Robert Gordon's paper '*Is US economic growth over?*' got a lot of attention. He argues:

> 'There was virtually no growth before 1750, and thus there is no guarantee that growth will continue indefinitely. Rather, the paper suggests that the rapid progress made over the past 250 years could well turn out to be a unique episode in human history...

> 'The analysis links periods of slow and rapid growth to the timing of the three industrial revolutions (IR's), that is, IR #1 (steam, railroads) from 1750 to 1830; IR #2 (electricity, internal combustion engine, running water, indoor toilets, communications, entertainment, chemicals, petroleum) from 1870 to 1900; and IR #3 (computers, the web, mobile phones) from 1960 to present. It provides evidence that IR #2 was more important than the others and was largely responsible for 80 years of relatively rapid productivity growth between 1890 and 1972. Once the spin-off inventions from IR #2 (airplanes, air conditioning, interstate highways) had run their course, productivity growth during 1972-96 was much slower than before.

> 'In contrast, IR #3 created only a short-lived growth revival between 1996 and 2004. Many of the original and spin-off inventions of IR #2 could happen only once – urbanization, transportation speed, the freedom of females from the drudgery of carrying tons of water per year, and the role of central heating and air conditioning in achieving a year-round constant temperature.

> 'Even if innovation were to continue into the future at the rate of the two decades before 2007, the U.S. faces six headwinds that are in the process of dragging long-term growth to half or less of the 1.9 percent annual rate experienced between 1860 and 2007. These include demography, education, inequality, globalization, energy/environment, and the overhang of consumer and government debt... [F]uture growth in consumption per capita for the bottom 99 percent of the income distribution could fall below 0.5 percent per year for an extended period of decades.'
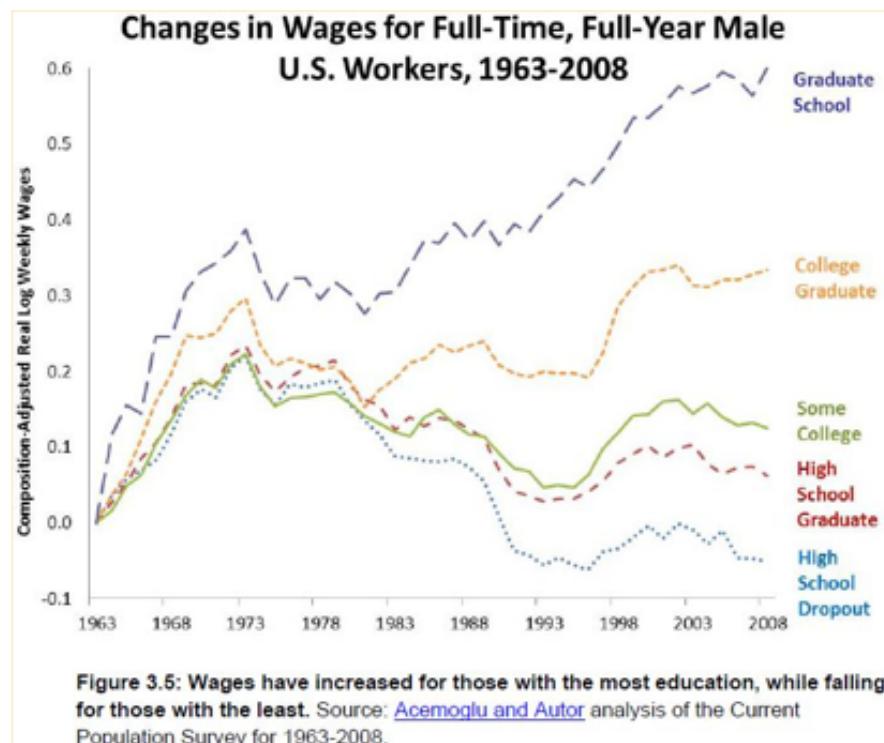
3. 'The end of work'. 'There has been no stagnation in technological progress or aggregate wealth creation as is sometimes claimed. Instead, the stagnation of median incomes primarily reflects a fundamental change in how the economy apportions income and wealth. *The median worker is losing the race against the machine*' (Brynjolfsson, emphasis added). Brynjolfsson's *Race Against the Machine* (2011) is an interesting rejection of Gordon's argument that the current computer-led IR#3 will not lead to changes similar to previous revolutions. However, he accepts that technological changes over the past thirty years have hurt employment and concentrated gains in the hands of the wealthy and best educated.

---

[256] However, while productivity declined from the 1960s level in the 1970s, 1980s, and 1990s, it rose again 2000-2010. There are obviously lots of problems with productivity figures. How do you measure greater choice? How do you measure the benefit of free things like Wikipedia or Google? '[M]ost government services are simply valued at cost, which implicitly assumes zero productivity growth... Health care productivity is poorly measured and often assumed to be stagnant, yet Americans live on average about 10 years longer today than than they did in 1960. This is enormously valuable, but it is not counted in our productivity data.' (Brynjolfsson)

High skills v low skills… Since the mid-1980s, the weekly wages of those without a high school diploma or with just a high school diploma have fallen while wages of those with the best education have risen.

Graph: weekly wages of groups with different education levels

Changes in Wages for Full-Time, Full-Year Male U.S. Workers, 1963-2008

Figure 3.5: Wages have increased for those with the most education, while falling for those with the least. Source: Acemoglu and Autor analysis of the Current Population Survey for 1963-2008.

The fact that the pay of the better educated has risen at the same time as the number of the better educated has also risen strongly suggests a rise in demand for the skills of the better educated.

Superstars v everyone else… The reason for the difference between rising GDP per capita and stagnant median incomes rests on the difference between *mean* and *median*: people like Bill Gates have a big effect on the mean, not on the median. Trillions of dollars of wealth have been created but *most of it has been concentrated in a relatively small number of hands*. Income has grown faster for the top 10% than for everybody else; faster for the top 1% than the rest of the top decile; faster for the top 0.1% than the top 1% and so on. The ratio of CEO pay to average worker pay increased from 70:1 in 1990 to 300:1 in 2005. In various ways, technology encourages 'winner takes all' markets in which a tiny number deemed 'the best' get most of the rewards (in sports and entertainment as well as technology). This can be economically efficient *and* socially damaging, argue some.

Capital v labour… Corporate profits as a share of GDP are at a 50-year high while compensation to labor is at a fifty year low. These figures probably understate the case since the labour figures include wages paid to superstars.

Many reject Gordon's argument regarding technology and think that the effects of the computer revolution will be at least as big as previous industrial revolutions, largely because of some of the things described above in this essay. Gordon asks: Imagine a choice between A) being able to keep everything from pre-2002 (including plumbing) but no electronics post-2002, or B) keep everything invented since 2002 but no plumbing. He recounts that audiences always prefer A. However, Kevin

Kelly [points out](#) that millions across the developing world *when faced with a real version of this choice have chosen B*. Kelly argues we are only in Year 20 of the 3rd Revolution (1992 kicked off the internet and computers before then changed little outside Government). Just as the 1st and 2nd IRs changed other things (like ideas of property) that bootstrapped higher growth, so the 3rd will change things we have not seen yet.

> 'In the year 2095 when economic grad students are asked to review this paper of Robert Gordon and write about why he was wrong back in 2012, they will say things like "Gordon missed the impact from the real inventions of this revolution: big data, ubiquitous mobile, quantified self, cheap AI, and personal work robots. All of these were far more consequential than stand alone computation, and yet all of them were embryonic and visible when he wrote his paper. He was looking backwards instead of forward."...'

However, even optimists usually concede that, first, increasingly rapid technological progress will bring greater benefits to the rich and powerful, and, second, that it may, after three hundred years of failed Luddite predictions, have a significant impact on employment: just as the internal combustion engine led to a large reduction in the employment of horses (it became uneconomic to pay for their feed rather than use a machine), *might humans become uneconomic as automation spreads?*

Even around 2004, there were predictions in mainstream economics that certain jobs such as driving, legal analysis, or journalism could not be automated, at least within decades. Even as this was being written, DARPA's Grand Challenge for autonomous vehicles was at work in the Mohave desert. In 2010, Google announced its driverless cars were a success. Blackstone Discovery in 2011 could analyse 1.5 million documents for less than $100,000, more accurately than many humans. StatsMonkey can write baseball news reports indistinguishable from bog standard online copy. Ironically, it may be that jobs requiring middling intellectual skills are decimated while jobs such as hairdressers and gardeners prove more resilient. After all, as NASA pointed out in 1965, until we have efficient intelligent robots, humans still have many advantages: 'Man is the lowest-cost, 150-pound, nonlinear, all-purpose computer system which can be mass-produced by unskilled labor' (NASA, 1965).

*

Better models: economics and epidemics

'Both when they are right and when they are wrong, the ideas of economists and political philosophers are more powerful than is commonly understood. Indeed the world is ruled by little else.' Keynes.

'I consider myself a proud neoclassicist... What I mean is that I prefer, when I can, to make sense of the world using models in which individuals maximize and the interaction of these individuals can be summarized by some concept of equilibrium. The reason I like that kind of model is not that I believe it to be literally true, but that I am intensely aware of the power of maximization-and-equilibrium to organize one's thinking - and I have seen the propensity of those who try to do economics without those organizing devices to produce sheer nonsense when they imagine they are freeing themselves from some confining orthodoxy.' Krugman, 1996.

'[T]he mathematics that the social scientists employ and the mathematical physics they use as their model are the mathematics and mathematical physics of 1850.' Wiener, 1964.

From the 1970's, and particularly in places such as the Santa Fe Institute from the 1980's, different groups came together to discuss common themes concerning complex systems, combining physicists (e.g. those studying the then new field of 'chaos'), biologists (e.g. those studying evolutionary game theory), computer scientists (e.g. those studying genetic algorithms, von Neumann's 'cellular automata'), economists (e.g. Brian Arthur), and others. Ideas floated between fields: e.g. the papers that developed 'nondeterministic chaos' in the 1970's (cf. Endnote), Trivers' *The Evolution of Reciprocal Altruism* (1971), Holland's *Adaptation in Natural and Artificial Systems* (1975), Dawkins' *The Selfish Gene* (1976), John Maynard Smith's *Evolution and the Theory of Games* (1982), Axelrod's & Hamilton's *The Evolution of Cooperation* (later a very influential 1984 book by Axelrod partly about computer game theory tournaments).[257] Traditional economics was a particular target and physicists began developing '*econophysics*' - the application of modern physics methods to markets. (Cf. Waldrop's *Complexity* for an entertaining history of the SFI.) An Endnote sketches some of the history of economic thinking including neoclassical economics and the intellectual origins of the DSGE models described below. Here I will just touch on some new ideas linking ideas from various fields with economics, finance, and economic forecasting, particularly the ideas of some 'econophysicists'.

Dealing with plasma physics is simple in one way - we now know a huge amount about fundamental particle physics. However, it turns out that if you compress and heat plasmas, there is an enormous range of complex behaviour that constantly surprises us (many reference books are filled with the details). Given how much better we understand the constituent parts of plasma physics (atoms, protons, quarks etc) than we do humans, and given we have quantitative predictive models for the former and none for the latter, *why would people expect such simple models as used in traditional economics to tell us much about interesting problems?* (Buchanan, 2013)

After the perceived failure of Keynesian economics in the 1970s, Lucas and others developed the 'rational expectations' model.

> 'The Keynesian predictions suggested that inflation could pull society out of a recession; that, as rising prices had historically stimulated supply, producers would respond to the rising prices seen under inflation by increasing production and hiring more workers. But when US policymakers increased the money supply in an attempt to stimulate employment, it didn't work — they ended up with both high inflation and high unemployment, a miserable state called 'stagflation'. Robert Lucas and others argued in 1976 that Keynesian models had failed because they neglected the power of human learning and adaptation. Firms and workers learned that inflation is just inflation, and is not the same as a real rise in prices relative to wages…' (Farmer, Nature Vol 460, 6/8/09)

The 'rational expectations' model is based on assuming, first, that humans have *perfect information* and, second, that they *adapt instantly and rationally* to new information in order to maximise their gains. From 'rational expectations' grew 'dynamic stochastic general equilibrium' (DSGE) models.

Governments use two types of model: 1) 'econometric: empirical statistical models that are fitted to past data. These successfully forecast a few quarters ahead as long as things stay more or less the same, but fail in the face of great change.' 2) DSGE models which 'by their very nature rule out crises of the type we are experiencing now' (Farmer).

---

[257] The SFI discussions of the 1980's often led people back to an earlier set of books: Schrödinger's *What is Life* (1944), von Neumann's papers (1945-57) on computers, automata, and the brain (cf. Endnote), Wiener's *Cybernetics* (1948), Shannon's *The Mathematical Theory of Communication* (1949).

Farmer describes some of the problems with DSGE models. First, people do not behave 'rationally' but are driven by emotions such as fear, so DSGE models do not incorporate the behaviour that leads to panic and crisis. Second:

> 'Even if [the 'rational expectations' model were] a reasonable model of human behaviour, the mathematical machinery is cumbersome and requires drastic simplifications to get tractable results. The equilibrium models that were developed, such as those used by the US Federal Reserve, by necessity stripped away most of the structure of a real economy. There are no banks or derivatives, much less sub-prime mortgages or credit default swaps — these introduce too much nonlinearity and complexity for equilibrium methods to handle.'

Also, these models assume that 'shocks' come from 'outside' the market, such as in the form of news. However, while some big moves are clearly related to big news (e.g. outbreak of war, Eisenhower's heart attack, 9/11), others occur without such news (e.g. 3/9/1946, 4/6/1962, October 1987, 6/5/2010 'the flash crash'). Various studies have shown that *most* of the big movements are *not* connected to big news - *they are self-generated by market dynamics* - and in general prices are more volatile than standard models predict. As algorithmic trading takes over, this seems to be happening more often - Nanex reports many fluctuations which are inexplicable in terms of markets responding to new information in the real world (i.e signals outside the HFT system itself). (Buchanan, 2013, and cf. Endnote for more on the origin and problems of these models.)

The problems with these models mean that many decisions by policy-makers, regulators and others are taken in a very unscientific manner.[258] Farmer describes how 'agent-based models' (ABMs) are different and might improve our understanding. In an ABM, every individual agent is represented as a 'distinct software individual'.

Robert Axelrod, one of the pioneers of ABMs, wrote:

> 'Agent-based modelling is a third way of doing science. Like deduction, it starts with a set of explicit assumptions. But unlike deduction, it does not prove theorems. Instead, an agent-based model generates simulated data that can be analyzed inductively. Unlike typical induction, however, the simulated data come from a rigorously specified set of rules rather than direct measurement of the real world. Whereas the purpose of induction is to find patterns in data and that of deduction is to find consequences of assumptions, the purpose of agent-based modelling is to aid intuition. Agent-based modelling is a way of doing thought experiments…
>
> 'The simulation is necessary because the interactions of adaptive agents typically lead to nonlinear effects that are not amenable to the deductive tools of formal mathematics.'

Agents can be programmed to operate with different models that can include fear, uncertainty, greed, and error, while behavioural information can be taken from sources such as surveys and cognitive science. All their interactions can be tracked. The models can be run thousands of times to gather robust statistical data.

> '[ABMs] do not rely on the assumption that the economy will move towards a predetermined equilibrium state, as other models do. Instead, at any given time, each agent acts according to its current situation, the state of the world around it and the rules governing its behaviour.'

---

[258] 'When the SEC changes trading rules, it typically has either flimsy or modest support from econometric evidence for the action, or else no empirical evidence and the change is driven by ideology.' Rob Axtell.

ABMs have already yielded useful help for policy-makers in some areas, such as epidemiology. Epstein describes how ABMs can help deal with epidemics:

> 'Classical epidemic modelling, which began in the 1920s, was built on differential equations. These models assume that the population is perfectly mixed, with people moving from the susceptible pool, to the infected one, to the recovered (or dead) one. Within these pools, everyone is identical, and no one adapts their behaviour. A triumph of parsimony, this approach revealed the threshold nature of epidemics and explained 'herd immunity', where the immunity of a subpopulation can stifle outbreaks, protecting the entire herd. But such models are ill-suited to capturing complex social networks and the direct contacts between individuals, who adapt their behaviours - perhaps irrationally - based on disease prevalence...'

Now 'the cutting edge in performance' for an ABM epidemic model is the Global-Scale Agent Model (GSAM)2.

> 'This includes 6.5 billion distinct agents, with movement and day-to-day local interactions modelled as available data allow. The epidemic plays out on a planetary map, colour-coded for the disease state of people in different regions — black for susceptible, red for infected, and blue for dead or recovered... For the United States, the GSAM contains 300 million cyber-people and every hospital and staffed bed in the country...

> 'In the wake of the 11 September terrorist attacks and anthrax attacks in 2001, ABMs played a similar part in designing containment strategies for smallpox...' (Epstein)

Running ABM simulations are not 'predictions'.

> '... [They can provide] a 'base case' which by design is highly unrealistic, ignoring pharmaceuticals, quarantines, school closures and behavioural adaptations. It is nonetheless essential because, base case in hand, we can rerun the model to investigate the questions that health agencies face. What is the best way to allocate limited supplies of vaccine or antiviral drugs? How effective are school or work closures?'

> 'Because they are rule-based, user-friendly and highly visual, they are natural tools for participatory modelling by teams - clinicians, public-health experts and modellers. The GSAM executes an entire US run in around ten minutes, fast enough for epidemic 'war games', giving decision-makers quick feedback on how interventions may play out. This speed may even permit the real-time streaming of surveillance data for disease tracking, akin to hurricane tracking.' (Epstein)

ABMs may also help in economics.

> 'Rob Axtell ... has devised firm dynamics models that simulate how companies grow and decline as workers move between them. These replicate the power-law distribution of

company size that one sees in real life: a very few large firms, and a vast number of very small ones with only one or two employees.' (Farmer)[259]

Farmer *et al* built an ABM to explore how leverage (the ratio of total assets owned to the wealth of the borrower) affects prices:

'There are four types of agents in this model. 'Noise traders', who trade more or less at random, but are slightly biased toward driving prices towards a fundamental value; hedge funds, which hold a stock when it is under-priced and otherwise hold cash; investors who decide whether to invest in a hedge fund; and a bank that can lend money to the hedge funds, allowing them to buy more stock...

'This agent-based model shows how the behaviour of the hedge funds amplifies price fluctuations, and in extreme cases causes crashes. The price statistics from this model look very much like reality. It shows that the standard ways banks attempt to reduce their own risk can create more risk for the whole system.'[260]

Specifying how agents decide and the rules for interaction has been done mainly by commonsense and guesswork. In the future 'we must proceed systematically, avoiding arbitrary assumptions, carefully grounding and testing each piece of the model against reality and introducing additional complexity only when it is needed.'

'Done right, the agent-based method can provide an unprecedented understanding of the emergent properties of interacting parts in complex circumstances where intuition fails. A thorough attempt to understand the whole economy through agent-based modelling will require integrating models of financial interactions with those of industrial production, real estate, government spending, taxes, business investment, foreign trade and investment, and with consumer behaviour. The resulting simulation could be used to evaluate the effectiveness of different approaches to economic stimulus, such as tax reductions versus public spending…

'In principle it might even be possible to create an agent-based economic model capable of making useful forecasts of the real economy, although this is ambitious. Creating a carefully crafted agent-based model of the whole economy is, like climate modelling, a huge undertaking. It requires close feedback between simulation, testing, data collection and the development of theory. This demands serious computing power and multidisciplinary collaboration among economists, computer scientists, psychologists, biologists and physical scientists with experience in largescale modelling. A few million dollars - much less than 0.001% of the US financial stimulus package against the recession - would allow a serious start on such an effort.'

---

[259] Other regulators are experimenting with ABMs. For example, they have paid attention to an open-source ABM known as the AMES Wholesale Power Market Test Bed used for testing the design of energy markets. Before NASDAQ switched to a decimal, rather than fractional, price system, they hired BiosGroup (Santa Fe) to do an ABM, which shaped how they made the change, because, said Mike Brown, 'Over ten years on the NASDAQ Board, I grew increasingly disappointed in our approach to studying the consequences of proposed market regulations, and wanted to try something different.' Companies are also using ABMs. E.g. Procter & Gamble and South West airlines used ABMs to investigate logistics and routing decisions.

[260] Blake Lebaron has developed Brian Arthur's asset price models into large scale agent-based simulations. There are 'trader' agents who try to predict the future, some programmed to look at 'fundamentals', and others to be more like 'speculators'. These models simulate the patterns seen in real markets, such as crashes and clustered volatility. (Cf. *'Agent-Based Financial Markets: Matching Stylized Facts with Style,'* 2010) and his website: http://people.brandeis.edu/~blebaron/wps.html

Developments with such models are connected to the analysis of *network topology and dynamics* described above. The network of *economic transactions, like the network of a cell or internet servers, has a scale-free topology (*i.e. hub and spoke) *which makes it vulnerable to crisis*. For example, only a small number of large banks (large hubs) trade with a large number of other smaller banks, making the system vulnerable to a crisis in a single large bank, and international trade and EU foreign direct investment follow a similar pattern (i.e. dominated by a small number of very big firms).[261] The next generation of analysis should 'begin to merge the description of individual agents strategies with their coevolving networks of interactions. We should then be able to predict and propose economic policies that favor networks structures that are more robust to economic shocks and that can facilitate integration or trade' (Schweitzer et al, *Science* 2009). Such tools will need to cope with *heterogeneous agents* and *positive feedback*.

Our unprecedented abilities to use smartphones (with GPS, accelerometers, microphones, cameras, and other sensors) and the internet mean that we are also now able to use society itself as an experimental data set to further knowledge about monitoring, quantitative modelling, and predictions for complex systems. Already, physicists are analysing the vastly complex physical and social networks of cities.

> 'Analogously to what happened in physics, we are finally in the position to move from the analysis of the "social atom" ... to the quantitative analysis of social aggregate states, as envisioned by social scientists at the beginning of the past century... [I.e.] large-scale social systems consisting of millions of individuals that can be characterized in space (geographic and social) and time. The shift from the study of a small number of elements to the study of the behavior of large-scale aggregates is equivalent to the shift from atomic and molecular physics to the physics of matter. The understanding of how the same elements assembled in large number can give rise ... to different macroscopic and dynamical behaviors opens the path to quantitative computational approaches and forecasting power. Yet at the same time, the study of social aggregate states present us with all the challenges already faced in the physics of matter, from turbulence to multiscale behavior... Complex systems and network theory, mathematical biology, statistics, nonequilibrium statistical physics, and computer science all play a key role' (Vespignani).[262]

Such tools are also being developed by intelligence services, militaries and associated consultancies.[263]

---

[261] The biggest five US banks account for >½ of all bank assets in the US (up from 29% two decades ago) and 95% of the top derivatives markets are accounted for by five banks (JP Morgan's notional gross derivatives exposure is >\$75 trillion (net 'a couple of trillion'). G. Tett, FT, 14/5/2012.

[262] Cf. Buchanan, Farmer, & Epstein pieces on 'agent-based modelling' (Nature) and the Science special issue.

[263] Palantir Technologies has used network analysis to search for IEDs and terrorists. Krebs modeled the 9/11 gang and found that Atta stood out for Degree (intensity of activity in the network), Closeness (ability to access others in the network), and Betweenness (control over the flow in the network). Carley has developed the 'Organizational Risk Analyzer' (ORA) to provide more dynamic analysis.

Table: a summary of differences between traditional models and ABMs (Miller & Page).

Traditional tools	ABM Models
Precise	Flexible
Little process	Process-oriented
Timeless	Timely
Optimising	Adaptive
Static	Dynamic
1,2, or infinite agents	1,2,...n agents
Vacuous	Spacey/networked
Homogenous	Heterogeneous

Table: paraphrasing Beinhocker on the main differences between 'traditional' and 'complexity' economics

	Complexity Economics	Traditional Economics
Dynamics	Open, dynamic, nonlinear systems in dis-equilibrium	Closed, static, linear systems in equilibrium
Agents	Modelled individually, inductive heuristics, cognitive biases, learn and adapt	Modelled collectively, deductive calculations, complete information, no errors or biases, no need for learning or adaptation
Networks	Model interactions, networks change over time	Treat agents as interacting indirectly through market mechanisms like auctions
Emergence	No distinction between micro and macro, macro patterns emerge from micro	Micro and macro are separate disciplines
Evolution	Differentiation, selection, amplification provides innovation and generates growth and complexity	No mechanism for endogenous growth, innovation, and complexity

*

Some conclusions…

Our fragile civilisation is vulnerable to large shocks such as the Great Depression, or death on the million or billion scale by accident (e.g pandemics or asteroids) or design (precise biological weapons perhaps targeted at an ethnic group *via* genetic markers as we develop the promise and peril of biotechnology).[264] A continuation of traditional human politics as it was during 6 million years of hominid evolution – an attempt to secure in-group cohesion, prosperity and strength in order to dominate or destroy nearby out-groups in competition for scarce resources - using advanced technologies could kill billions.[265]

World War II began with cavalry charges - an ancient technology and one well understood by senior politicians. It ended with nuclear weapons and computers - technologies understood by few among ruling elites who now have to contend with even greater complexity. Market, scientific, and technological trends described above are having various effects:

1) The success of markets and science in bringing in-group prosperity and power for states has made it hard to resist what Gibbon called 'creative emulation': military applications of science mean that any in-group that rejects its exploitation is vulnerable to defeat or destruction.
2) They are bringing nearly all humans into 'informational proximity', for good and ill.
3) *'Revolution must necessarily begin with atheism' (Dostoyevsky).* Markets and science undermine ancient religions, traditional morals and social structures. For example, they require, allow, encourage and protect a private sphere which can be used to challenge all authority, such as strict controls of women. Unsurprisingly, support for female education (so important for the world) provokes hate and reaction. The weakening of evolved social norms regarding sex (not outside marriage) and work (the able-bodied should support themselves) have extended human freedom *and* caused major social problems. Both required the collapse of elite moral authority.
4) Science has undermined traditional ideas about a religious basis for ethics but has not produced 'a rational basis for morality', the hope of Kant and others.[266]

---

[264] The global population of people with IQ +4 SD (>160) is ~250k. ~1% of the population are psychopaths so there are perhaps ~2-3,000 with IQ ≈ Nobel/Fields winner. The psychopathic +3SD IQ (>145; average science PhD ~130) population is 30 times bigger. A subset will also be practically competent. Some of them may think, '*Flectere si nequeo superos, / Acheronta movebo*' ('If Heav'n thou can'st not bend, Hell thou shalt move', the Aeneid). *Board et al (2005):* high-level executives are *more* likely than inmates of Broadmoor to have one of three personality disorders (PDs): histrionic PD, narcissistic PD, and obsessive-compulsive PD. *Mullins-Sweatt et al (2010):* successful psychopaths are *more* conscientious than the unsuccessful. While most of those at the cutting edge of technology are, perhaps inevitably, optimists, Bill Joy, a Silicon Valley legend, gave a famous warning in his 2000 essay that traditional pursuit of technological progress risks various catastrophes ('*Why the future doesn't need us*', Wired, 2000; NB. quote marks are missing in this version).

[265] All societies divide people into in/out group, child/adult, kin/nonkin, married/single, good/evil, sacred/profane. The boundaries are fortified with ritual and taboo. The territorial instinct arises 'when some vital resource serves as a density-dependent factor'; ie. 'the growth of population density is slowed incrementally by an increasing shortage of food, water, nest sites, or the entire local terrain available to individuals searching for these resources.' Death rates and birth rates come into balance and territorial behaviour tends to evolve because individuals predisposed to defend private resources for themselves and their social group pass on more genes. 'In contrast, the growth of other species is not leveled off by limiting resources but by rising amounts of emigration, disease, or predation. When such alternative density-dependent factors are paramount, and resource control is therefore not required, territorial defense usually does not evolve as a hereditary response… Humanity is decidedly a territorial species.' Unlike social insects mammals are concerned with themselves and close kin, 'they resist committing their bodies and services to the common good', hence 'nonhuman mammalian species are far less organized than the insect societies. They depend on a combination of dominance hierarchies, rapidly shifting alliances, and blood ties. Human beings have loosened this constraint and improved social organization by extending kinship ties to others through long-term contracts.' Cheating 'serves as the principal source of hostile gossip and moralistic aggression by which the integrity of the political economy is maintained.' (E.O. Wilson) One of the world's most prestigious scientific journals has recently run a special issue on 'human conflict' (an example of consilience in action).

[266] 'Those who have abandoned God cling that much more firmly to morality.' Nietzsche.

5) Many of the things mentioned in this essay (such as encrypted 'darknets', personal fabrication, drones) not only are transforming the economy but also give great power to individuals and small groups (as well as traditional militaries and intelligence agencies) to conduct operations in the physical, virtual or psychological worlds (destruction, 'information operations',[267] cyberwar etc) *potentially without detection*.[268] A.Q. Khan has spread nuclear technology far and wide and who knows where the Soviet biowar scientists went after 1991 and what they have done since?

6) They are evolving faster than the ability of bureaucracies to adapt, bringing a perpetual sense of crisis and fragility that undermines support for markets and science in a vicious feedback loop. Many new technologies are blamed for the loss of jobs.[269] Many technologies are opposed with ignorant and/or dishonest arguments (e.g. GM foods). Financial markets - plagued by huge debts, complex mis-regulation, and a widespread culture of cheating customers - are now dominated by high-frequency 'algorithmic trading' involving vast volumes traded in fractions of a second, far beyond the monitoring or understanding of politicians and regulators. There is competition for computer trading bases in specific locations based on calculations of Special Relativity as the speed of light becomes a factor in minimising trade delays (cf. *Relativistic statistical arbitrage*, Wissner-Gross, *Physical Review E*, 2010). Mini 'flash crashes' blow up and die out faster than humans can notice as robots fight each other in financial markets. *There seems scope for worse disasters than 2008 which would further damage the moral credibility of decentralised markets.*[270]

7) The protection of private space, and the implicit or explicit acceptance of public discussion about the wisdom of different ideas, creates pressures for open discussion of other public affairs that the (non-democratic) state may consider threatening. E.g. China is trying to balance sufficient *openness* for science and business against sufficient *control* to prevent Tiananmen-style protests: could it lead in robotics (as it plans) *and* seriously limit scientists' access to the internet? Private property, and the broader concept of 'equal under the law' that its protection has encouraged in England and America, are direct challenges to the authority and power of the state, hence the lack of such concepts in places like Russia has shaped their political history and their recent attempts to 'modernise' have collided not just with deep-rooted culture but with ruling elites's fears (justified in one sense) that real markets will undermine their control. While, democracy is demonstrably not *necessary* for thriving markets, it seems to be a likely consequence of developed market institutions: 'equal votes' is a more promising slogan in a culture that is already used to 'equal protection for property'.

8) While the worst-educated often feel markets are an exploitative racket, the best educated often feel alienation from (or hatred for) market-based societies which they feel do not value them, are repulsed by what they think of as the mundanity and ugliness of commerce and the 'irrationality' of markets, and are embarrassed by what they think of as the parochialism of national governments. Many are educated sufficiently to feel resentment against existing order but insufficiently to create useful things or be responsible for others, like various characters in Dostoyevsky, and they can be extremely susceptible to intellectual fads such as hubris about

---

[267] Hezbollah's communications are more sophisticated than many traditional western political parties.

[268] There have been various attempts to describe how some of these trends might affect conflict: e.g. Cebrowski's 'Network Centric Warfare', van Creveld's '*Transformation of War*', Lind's '*Fourth Generation Warfare*', Robb's swarms of guerrillas using the tools and techniques of Silicon Valley ('open source war'). The Pentagon's latest attempt is the mostly classified 'Air-Sea Battle Concept' (a 2013 unclassified summary here).
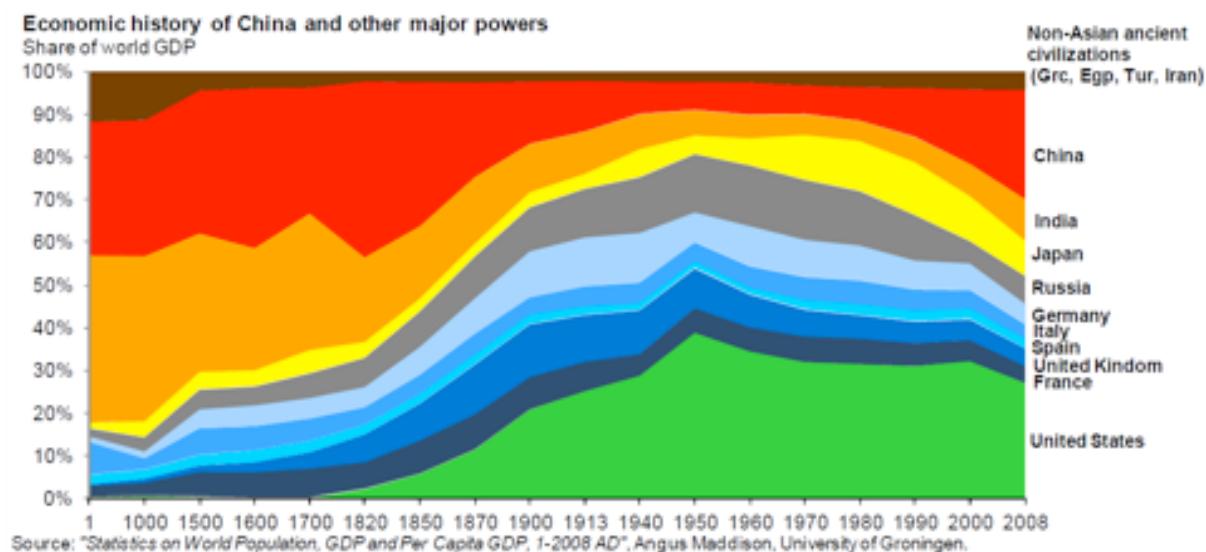
[269] Will improving technology be seen as good by most people? 'At the height of its power, ... Kodak employed more than 14,000 people and was worth $28 billion. They even invented the first digital camera. But today Kodak is bankrupt, and the new face of digital photography has become Instagram. When Instagram was sold to Facebook for a billion dollars in 2012, it employed only 13 people. Where did all those jobs disappear?' Lehrer. For a different view, cf. *Race Against the Machines*, 2012.

[270] 'The Flash Crash' of 9 May 2010 saw the Dow lose hundreds of points in minutes. We urgently need firewalls between low-risk financial institutions guaranteed by taxpayers and high-risk institutions that are allowed to go bust and wipe out share and bond-holders. Cf. This interesting talk by Robert May on modern finance, an example of trans-disciplinary study, and this paper by Cliff & Northrop for the UK Government on the financial system.

their ability to engineer complex systems.[271] It seems likely that riches will continue to concentrate in the hands of a few, making these problems worse.

9) The global population is growing towards ~10^{10}. Most of the extra 2-3 billion that will come in the next few decades will be in cities, which also increases the energy per capita consumed. The scale of global urbanisation is now equivalent to building five Londons per year. The world's urban population will rise from roughly 50 percent in 2010 to about 60 percent (~5b) by 2030. Urban centers are estimated to generate 80 percent of economic growth and urban construction 'over the next forty years could roughly equal the entire volume of such construction to date in world history.' (The National Intelligence Council's '*Global Trends 2030*'.) The combination of accelerating urbanisation and huge growth in livestock density also has implications for disease.

10) As well as growing, most of the world is ageing. (China's working age population has *already* peaked, it will soon have to deal with tens of millions more old people while the labour force shrinks, and it faces growing old *before* most people become middle-class.) While this puts finances under strain, it also reduces pressure in the developing world as 'youth bulges' are correlated with violence. This growing population is also increasingly mobile.

11) Economic power is shifting rapidly. It took Britain about 150 years (18th-19th Centuries) to double GDP; the USA and Germany took 30-60 years; China and India are doing it with 100 times more people than Britain did in a tenth of the time. The developing world already produces more than half of global economic growth. By 2025, China is projected to produce about a third of global growth and will soon pass the $15,000 per capita threshold (which has often been a trigger for democratisation movements). The developing world will have hundreds of millions more middle class consumers with some discretionary income.[272] This demand will force huge economic change and create, change, and destroy companies.
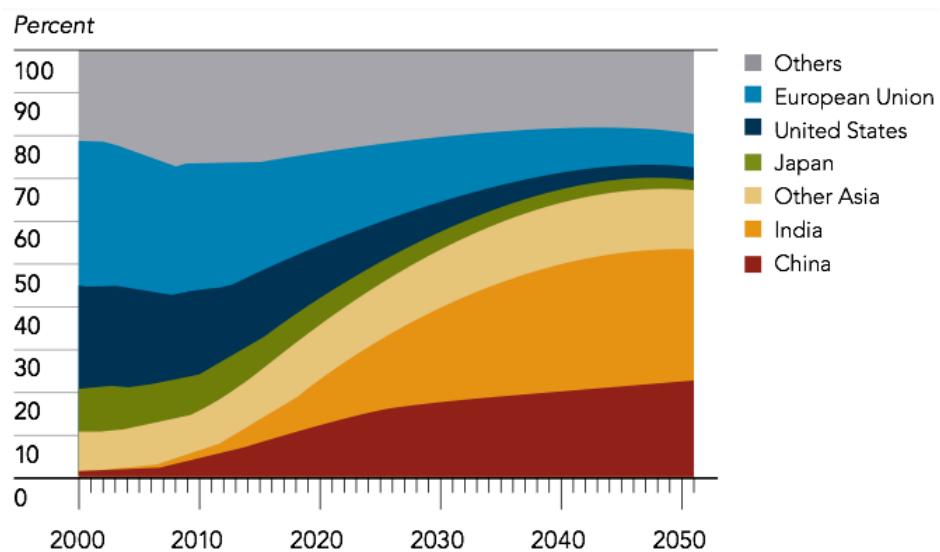
Different countries' changing share of world GDP

Economic history of China and other major powers
Share of world GDP
Source: "Statistics on World Population, GDP and Per Capita GDP, 1-2008 AD", Angus Maddison, University of Groningen.

---

[271] 'At least they take me seriously,' Brecht is supposed to have said about his support for the Soviet Union.

[272] Goldman Sachs predicts that the 'next eleven' (Bangladesh, Egypt, Indonesia, Iran, Mexico, Nigeria, Pakistan, the Philippines, South Korea, Turkey, Vietnam) will overtake the EU-27 by 2030. The developing world's demand for infrastructure, consumer goods, and new plants and equipment will raise global investment; if global savings do not rise, then there will be upward pressure on interest rates.

Share of global middle class consumption (Source: NIC, 2013)

12) As the world moves towards 10 billion people and more consuming a Manhattan-style 10^4 watts per capita, competition for resources will grow. 'Demand for food, water, and energy will grow by approximately 35, 40, and 50 percent [by 2030]... Nearly half the world's population will live in ares experiencing severe water stress' (NIC, 2013). Big changes in energy policy, such as shale gas providing US energy independence, will have big political effects (e.g what will change in the Middle East if America no longer sees its oil as a vital interest?). While there are big problems, we can also see the contours of a new landscape. There is easily enough energy, water, land, and minerals for billions more than currently exist to live lives like the richer populations of the West *if* we continue to innovate and become more productive. Our scientific, technological and economic methods have inevitably been relatively crude therefore our consumption and environmental impact have put us on an unsustainable trajectory. We can see how more precise understanding and engineering, with better institutions (such as markets structured to protect the commons), could enable us to escape trends that threaten disaster.[273]

13) Asia is militarising as it gets richer. On one hand, China's power is growing rapidly. On the other, there is no reason to think that China will avoid similar major crises that other rising powers have suffered, such as a major financial crisis causing a Depression. Imagine a combination of a financial crisis, depression, and a declaration of independence by Taiwan: might a future Thucydides write, 'the avowed cause of the war was Taiwanese "freedom" but the real cause was America's fear of China's growing power' (or 'China's fear of collapse')?

14) In Europe, a dysfunctional EU has members facing rapidly ageing populations, welfare and pension systems that are insolvent, health systems that will require huge new resources, large long-term and youth unemployment, large and rising debts, anti-entrepreneurial political cultures, low growth, and a dysfunctional single currency and political institutions. Those who control EU institutions have repeatedly used every crisis as a spur to further political integration and this approach will probably have at least 'one more heave'. Britain remains trapped, neither

---

[273] Many environmental arguments rely on the 'IPAT' equation. This claims that environmental Impact (I) is equal to Population (P) times Affluence (A) times Technology use (T): (I=PxAxT). The assumption that pollution necessarily increases as the economy grows is wrong. We have already proved that applied science plus markets can cut pollution substantially while the economy grows (e.g with CFCs and acid rain). The US spends a trillion dollars per year on energy consumption but only $5 billion on energy R&D (200:1); it spends 2.6 trillion dollars on healthcare but only $30 billion on National Institutes of Health (100:1). Amid huge waste, we skimp on the important.

committed to the project nor daring to leave and pursue a different strategy *pour encourager les autres*.

15) Democracies are notoriously unable to cope with long-term strategies and commitments: consider Pericles' warning to the Athenians to stick to a long-term plan and the demos' failure to do so. Dictatorships are notoriously warlike. Various issues explored in this essay could encourage us to think of ourselves as part of a common family exploring our place in a world of vast resources, rather than as hostile groups competing for scarce resources on a crowded planet. They also require international cooperation and limits on what individual countries are allowed to do by others but international and supranational institutions are often scaled-up versions of national failures, minus accountability (power is exercised supranationally but is accountable nationally).[274] Neither the EU nor UN are remotely fit for what we need. The change in economic power is already leading to understandable demands from the developing world for changes to international organisations but the West has no new vision. On one hand, people in the West sometimes demand 'something must be done' about suffering abroad; on the other hand, people shrink from the consequences, inevitably messy, of 'savage wars of peace' and 'The blame of those ye better, The hate of those ye guard'. 'Between the idea / And the reality... Falls the shadow' and our politicians seldom analyse, clearly and publicly, such problems in the way Athenian politicians analysed Melos or Miletus. Imagine a Prime Minister saying on TV, 'Given I do not think the possible costs of removing X are worthwhile, it would be a foolish and counterproductive gesture to pretend we are serious by a token attack on X, so I will continue to focus my efforts elsewhere and not pretend we are going to solve the problem of X.'

While our ancestor chiefs understood bows, horses, and agriculture, our contemporary chiefs (and those in the media responsible for scrutiny of decisions) generally do not understand their equivalents, and are often less experienced in managing complex organisations than their predecessors, while dangers are greater.

<div align="center">*</div>

Political philosophy.

'Our morals are not the conclusions of our reason.' Hume.

'Revolution must necessarily begin with atheism.' Dostoyevsky.

'The madman jumped into their midst and pierced them with his eyes. "Whither is God?" he cried. "I will tell you. We have killed him, you and I. All of us are his murderers. But how did we do this? How could we drink up the sea? Who gave us the sponge to wipe away the whole horizon? What were we doing when we unchained this earth from its sun? Whither is it moving now? Whither are we moving?... Are we not straying as through an infinite nothing? Do we not feel the breath of empty space? Has it not become colder? Is not night continually closing in on us? Do we not need to light lanterns in the morning? Do we hear nothing as yet of the noise of the gravediggers who are burying God? Do we smell nothing as yet of the divine decomposition? Gods, too, decompose. God is dead. God remains dead. And we have killed him. How shall we comfort ourselves, the murderers of all murderers? What was holiest and mightiest of all that the world has yet owned has bled to death under our knives: who will wipe this blood off us? What water is there for us to clean ourselves? What festivals of atonement, what sacred games shall we have to invent? Is not the greatness of this deed too great for us? Must we ourselves not become gods simply to appear worthy of it? There has never been a greater deed; and whoever is born after us – for the sake of this deed he will belong to a higher history than all history hitherto."...

---

[274] Many supporters of the EU confuse 'power' / 'sovereignty' and cooperation / law.

'They [the English] have got rid of the Christian God and now feel obliged to cling all the more firmly to Christian morality… When one gives up Christian belief one thereby deprives oneself of the right to Christian morality. For the latter is absolutely not self-evident… Christianity is a system, a consistently thought out and complete view of things. If one breaks out of it a fundamental idea, the belief in God, one thereby breaks the whole thing to pieces: one has nothing of any consequence left in one's hands. Christianity presupposes that man does not know, cannot know what is good for him and what is evil: he believes in God, who alone knows. Christian morality is a command: its origin is transcendental; it is beyond all criticism, all right to criticise; it possesses truth only if God is truth – it stands or falls with the belief in God…

'Supposing all danger, the cause of fear, could be abolished, this morality [Christian, socialist] would herewith also be abolished: it would no longer be necessary, it would no longer regard itself as necessary! He who examines the conscience of the present-day European will have to extract from a thousand moral recesses and hiding-places always the same imperative, the imperative of herd timidity: "we wish that there will one day no longer be anything to fear!" One day – everywhere in Europe, the way and will to that day is called "progress".'..

'Behind all the moral and political foregrounds … a great physiological process is taking place and gathering greater and ever greater impetus - the process of the assimilation of all Europeans, … their growing independence from any definite milieu, which, through making the same demands for centuries, would like to inscribe itself on soul and body - that is to say, the slow emergence of an essentially supranational and nomadic type of man which … possesses as its typical distinction a maximum of the art and power of adaptation. This process of the becoming European, *the tempo of which can be retarded by great relapses but which will perhaps precisely through them gain in vehemence and depth - the still-raging storm and stress of "national feeling" belongs here, likewise the anarchism now emerging - this process will probably lead to results which its naïve propagators and panegyrists, the apostles of "modern ideas", would be least inclined to anticipate. The same novel conditions which will on average create a levelling and mediocritizing of man … are adapted in the highest degree to giving rise to exceptional men of the most dangerous and enticing quality. For … while the total impression produced by such future Europeans will probably be that of multifarious, garrulous, weak-willed and highly employable workers who need a master, a commander, as they need their daily bread; while, therefore, the democratisation of Europe will lead to the production of a type prepared for slavery in the subtlest sense: in individual and exceptional cases the strong man will be found to turn out stronger and richer than has perhaps ever happened before – thanks to the unprejudiced nature of his schooling, thanks to the tremendous multiplicity of practice, art and mask. What I mean to say is that the democratisation of Europe is at the same time an involuntary arrangement for the breeding of tyrants…' Nietzsche.*

Existing political philosophies are inadequate to serve as heuristics for decision-makers. Established political philosophies including traditional conservatism, liberalism and socialism, which form the basis of background heuristics for political leaders, cannot cope with evolutionary epistemology, either in biology or economics. Individuals among the traditional 'Left', 'Right', conservatives, neoconservatives, socialists, liberals and libertarians *may* accept some notion of man as an evolved creature but the politically active among all groups generally reject evolutionary biology as a basis for understanding human nature.

*Traditional Conservatism (*as per Oakeshott, for all his wisdom) has not coped with a world in which Religion (post-Darwin) and Nation (post-two world wars) ceased to be accepted conventions and became objects of intellectual scrutiny requiring intellectual defence. Since markets are inherently disruptive of established traditions, ideas, and values (always seeking new information for creative destruction), and are inevitably subject to occasional crises, traditional Conservatism is inherently at least ambivalent about, and often actively hostile to, markets. *Conservatism and liberalism (*and

'neoconservativism') try to accommodate the evolutionary nature of markets but are at best ambivalent about evolutionary biology. However, widespread hostility to markets as 'unjust' ('the rich cheat', 'the wrong people are rewarded'), hostility to free trade ('foreigners steal our jobs'), and hostility to 'irrational markets in which a lack of planning causes chaos and destruction', are unsurprising from an evolutionary perspective.

Socialism is in tune with some of our evolved instincts such as collectivist distrust of markets and free trade, and a desire for a 'wise benevolent chief' to command in the 'common interest' of the tribe. However, it struggles to cope with evolved instincts such as our preference to help closer relatives over unknown strangers ('kinship selection'), 'reciprocal altruism', or our instinct when under pressure to attack other out-groups rather than unite in class warfare. The survival of ethnic and nationalist hatreds, scepticism towards international or supranational institutions, and deep ambivalence towards welfare systems[275] are also unsurprising from an evolutionary perspective. Further, socialism suffered a major disaster with the observable fact of evolutionary markets' success versus experiments such as the Soviet Union or North Korea (though the disaster of the 2008 financial crisis has also discredited markets).

Broadly, those who self-define as 'Left' tend to apply evolutionary logic to God but deny it with markets, while those who self-define as 'Right' operate vice-versa. The Left and Right argue that *evolving information provides an explanation for complex emergent phenomena that appear to be designed*; the Left aims this argument at God, the Right at markets, and the former dislikes its implications for hubristic attempts to micromanage complex systems while the latter dislikes its implications for hubristic claims that 'God is dead'. Many of both tendencies see science as a sniper to be deployed at will rather than a coherent intellectual approach. Broadly, the Left fears that an evolved universal human nature undermines the idea of *equality*. The Right fears that it undermines the idea of *responsibility*: a world in which God is dead and morals are evolved, not transcendent, is, they fear, a world of nihilism and immorality. 'But what will become of men without God and immortal life? All things are lawful then, they can do what they like?', as Dimitri asks in *Brothers Karamazov*.

Traditional political philosophies have not been helped by modern 'social sciences'. Marx (economics), Freud (psychology), Durkheim (sociology), Boas (anthropology), Rawls (law) and others rejected the idea of their disciplines being based on an evolved universal human nature.[276] This profoundly shaped the 'the standard social science model' (SSSM), the basis of much modern economics, psychology, sociology and anthropology. In the early part of the 20th Century, Behaviourism reigned supreme. It was deemed 'unscientific' to consider mental phenomena for they were 'unobservable' (another consequence of Positivism). This is no longer tenable. Further, the spread of French 'post-modernism' like an anti-scientific virus through the humanities has been

---

[275] In-group solidarity encourages limited 'welfare', but it is hard to balance humane treatment and destructive incentives that undermine that solidarity.

[276] Boas (1858–1942): 'We must assume that all complex activities are socially determined, not hereditary.' Durkheim (1858 –1917): 'Every time that a social phenomenon is directly explained by a psychological phenomenon, we may be sure that the explanation is false… If we begin with the individual in seeking to explain phenomena, we shall be able to understand nothing of what takes place in the group… Individual natures are merely the indeterminate material that the social factor molds and transforms.' Unsurprisingly, many anthropologists and others attacked Napoleon Chagnon's famous work on the warlike hunter-gatherer tribe of the Yanomamo and spread false accusations about him.

disastrous as a few generations of students have been encouraged to think of science as 'just another perspective', a trend exposed in [the Sokal Hoax](#).[277]

Philosophy sacrificed its power as it ceased to be 'natural philosophy' and withdrew from other fields. For example, Wittgenstein was very influential in philosophy departments but his ideas about the foundation of knowledge in mathematics and logic was thoroughly rejected, and shown to be false, by mathematicians themselves (cf. Endnote).

> 'The fading of philosophy came to my attention in 1979, when I was involved in the planning of a conference to celebrate the hundredth birthday of Einstein… After acrimonious discussions, we agreed to have three committees… One committee was for scientists, one for historians of science, and one for philosophers of science. After the three committees had made their selections, we had three lists of names… With a few exceptions, I knew personally all the people on the science list. On the history list, I knew the names, but I did not know the people personally. On the philosophy list, I did not even know the names… Philosophers became insignificant when philosophy became a separate academic discipline, distinct from science and history and literature and religion… The latest masterpieces written by a philosopher were probably Friedrich Nietzsche's *Thus Spoke Zarathustra* in 1885 and *Beyond Good and Evil* in 1886.' Freeman Dyson, 2012.

Fukuyama's *The End of History and the Last Man* is often mocked as a naive and hubristic claim of victory for liberal democracy and triumph for Hegel's 'Last Man' - the 'man of the herd', as Nietzsche called him. Fukuyama's last chapter, however, admitted that the greatest philosophical attack on liberal democracy came from Nietzsche. Nietzsche mocked the Last Men who 'no longer have any idea what religions are supposed to be for and merely register their existence in the world with a kind of dumb amazement.' He mocked their desire for 'the universal green pasture of the herd, with security, safety, comfort and an easier life for all; their two most oft-recited doctrines and ditties are 'equality of rights' and 'sympathy for all that suffers' – and suffering itself they take for something that has to be *abolished*.' He mocked the modern, liberal enlightened scholar for the 'venerable, childlike and boundlessly stupid naivety there is in the scholar's belief in his superiority, in the good conscience of his tolerance, in the simple unsuspecting certainty with which his instinct treats the religious man as an inferior and lower type which he himself has grown beyond and *above* – he, the little presumptuous dwarf and man of the mob, the brisk and busy head- and handyman of 'ideas', of 'modern ideas'…' He mocked the way in which such people thought they could reject the Christian God but keep Christian morality, and he mocked their claims to be 'free spirits': 'The claim to independence, to free development, to *laisser aller*, is advanced most heatedly by precisely those for whom no curb *could be too strong* – this applies in politics, it applies in art.' He mocked and rejected the whole goal of modern liberal society: 'The collective degeneration of man down to that which the socialist dolts and blockheads today see as their 'man of the future' –

---

[277] Social science generally suffers from a lack of precise language, unscientific influence from ideology, tribalism, and folk psychology. Whereas the natural sciences have progressed by looking at the individual elements that make up the system, social sciences have rejected the idea of understanding the individual and have focused on 'the system'. The SSSM 'sees culture as an independent phenomenon irreducible to elements of biology and psychology, thus the product of environment and historical antecedents. In purest form [the SSSM] turns the intuitively obvious sequence of causation upside down: Human minds do not create culture but are themselves the product of culture. This reasoning is based … on the slighting or outright denial of a biologically based human nature.' (E.O.W) 'The more social the discipline, the more retarded its development,' says Trivers, and he is particularly hostile to 'cultural anthropology' and its 'bizarre mind states, such as believing that words have the power to dominate reality, that social constructs such as gender are much stronger than the 300 million years of genetic evolution that went into producing the two sexes – whose facts in any case they remain resolutely ignorant of, the better to develop a thoroughly word-based approach to the subject. In many ways, cultural anthropology is now all about self-deception' (Trivers). Cf. Nobel-winner Steve Weinberg's commentary in *Science and Its Cultural Adversaries* and his essay 'Against Philosophy'. If ethics are shaped only by 'culture' and cultures are diverse and 'equally moral', then what is the argument against slavery or torture?

as their ideal! – this degeneration and diminution of man to the perfect herd animal (or as they say, to the man of the 'free society'), this animalisation of man to the pygmy animal of equal rights and equal pretensions is *possible*, there is no doubt about that! He who has once thought this possibility through to the end knows one more kind of disgust than other men do – and perhaps also a new *task*!' When Zarathustra describes the Last Man, the crowd want to be turned into these Last Men; Zarathustra replied, 'For thus you speak: "Real are we wholly, and without faith or superstition." Thus you stick out your chests, but alas they are hollow.'

For Nietzsche, the spreading breakdown of belief in Christianity - an enormous event, beyond the comprehension of the people who had done it - also offered the promise of new philosophies that would restore the spirit of the pre-Socratic Greeks or the Renaissance, which he regarded as the true pinnacles of mankind:

> 'Individuals and generations can now fix their eyes on tasks of a vastness that would to earlier ages have seemed madness and a trifling with Heaven and Hell. We may experiment upon ourselves... Man is a rope fastened between animal and Superman...'

Modern philosophy has not known what to do with his challenge and as Dyson suggests Nietzsche is probably the last of the line of recognisable great philosophers stretching back to Thales of Miletus. A renaissance requires that it rebuild its connections to other subjects.

<p align="center">*</p>

After 1945, Dean Acheson quipped that Britain had failed to find a post-imperial role. It is suggested here that *this role should focus on making ourselves the leading country for education and science*: Pericles described Athens as 'the school of Greece', we could be the school of the world. Who knows what would happen to a political culture if a party embraced education and science as its defining mission and *therefore changed the nature of the people running it and the way they make decisions and priorities*. We already have a head start; we lack focus. Large improvements in education and training are easier to achieve than solving many other big problems and will contribute to their solution. Progress *could* encourage non-zerosum institutions and global cooperation - alternatives to traditional politics and destruction of competitors. However, the spread of knowledge and education is itself a danger and cannot eliminate gaps in wealth and power created partly by unequally distributed heritable characteristics.

The acceleration of scientific knowledge will continue to enhance the dangerous potential of our evolved nature and catastrophes are inevitable but better training and new institutions operating in the manner of an immune system will give us better chances to limit and survive them. It is hoped that this draft sketch will encourage others far more knowledgeable to develop an Odyssean curriculum and training system in a disciplined and open way.

Dominic Cummings
Version 2
25 August 2013

Endnotes

Origins of 'chaos theory'

Models, power laws, prediction.

Constant failure

Networks, modularity, and fractals

Some history of thinking about the limits of logic and computation, and 'P=NP?'

Quantum Computation

Prizes

Problems with Science

Intelligence, IQ, genetics, and extreme abilities

A scientific approach to science teaching

Gigerenzer: Bayes, risk, and education

Cognitive biases

Some sketches of the history of economic thinking

Boyd's OODA loop

Some interesting charts

An Odyssean Reading List

Endnote: Origins of 'chaos theory'

'[A principle of causation: "Like causes produce like effects."] This is only true when small variations in the initial circumstances produce only small variations in the final state of the system. In a great many physical phenomena this condition is satisfied; but there are other cases in which a small initial variation may produce a very great change in the final state of the system.' Maxwell

'The system of the 'universe as a whole' is such that quite small errors in the initial conditions can have an overwhelming effect at a later time. The displacement of a single electron by a billionth of a centimetre at one moment might make the difference between a man being killed by an avalanche a year later, or escaping.' Turing

> *'Embedded in the mud, glistening green and gold and black,*
> *was a butterfly, very beautiful and very dead.*
> *It fell to the floor, an exquisite thing, a small thing*
> *that could upset balances and knock down a line of*
> *small dominoes and then big dominoes and then*
> *gigantic dominoes, all down the years across Time.'*
> *Bradbury, 1952.*

Despite Poincaré's work in the 19th Century (see Section 1), the birth of 'chaos theory' is usually dated to 1961 when Edward Lorenz had a serendipitous discovery concerning equations he had constructed as a weather model.

Weather forecasting is of great interest to the military: for example, the mythical 'fog of Chulm' falsely spun by the Austrians as the cause of their defeat by Moltke in 1866, and the importance of forecasts to the timing and success of D-Day in 1944.[278] During World War I, an English physicist, Lewis Richardson, sat in his trench working on an idea to improve weather forecasting.

> 'One step at a time, he was constructing a cellular, numerical model of the weather over northwest Europe for a single six-hour period long past, treating the motions of the atmosphere as satisfying a system of differential equations that linked the conditions in adjacent cells from one interval to the next.' (Dyson)

This project combined a theory of eddy diffusion and a method 'to calculate approximate solutions to systems of differential equations resistant to the analytic approach.' After the war he published *'Weather Prediction by Numerical Process'* which, as Dyson writes, described a concept analogous with how modern computerised simulations work.

> 'Imagine a large hall like a theatre, except that the circles and galleries go right round through the space usually occupied by the stage. The walls of this chamber are painted to form a map of the globe… A myriad computers [human] are at work upon the weather of the part of the map where each sits, but each computer attends only to one equation or part of an equation. The work of each region is coordinated by an official of higher rank. Numerous little 'night signs' display the instantaneous values so that neighbouring computers can read them. Each number is thus displayed in three adjacent zones so as to maintain communication to the North and South on the map. From the floor of the pit a tall pillar rises to half the height of the hall. It carries a large pulpit on its top. In this sits the main

---

[278] The first weather map was printed in *The Times* on 1 April 1875. It was the idea of Darwin's cousin, Francis Galton.

charge of the whole theatre; he is surrounded by several assistants and messengers. One of his duties is to maintain a uniform speed of progress in all parts of the globe.' (Richardson)

Von Neumann unsurprisingly became interested in weather during the war and there are references to discussions concerning the atomic bombs, targeting, and weather forecasts.[279] As he developed the IAS computer program in 1945-6, he became convinced that meteorology ought to be its practical focus because the hydrodynamics of the atmosphere presented precisely the sort of nonlinear problem that computers could potentially revolutionise. He envisaged a *network* of computers across the globe sharing information on weather patterns. Over the next ten years, the entire field of meteorology was revolutionised.

Von Neumann was interested not only in weather forecasting but in weather *control*. He envisaged that weather control could significantly improve the environment in various parts of the earth and could also be used as a weapon (though he generally played down ideas about control and 'weaponising' in conversations with meteorologists). In 1950, the first computerised weather forecasts were performed. By 1952, several military agencies were developing similar programs with IAS's help. However, by 1955, von Neumann's other work, particularly with the Atomic Agency Commission, prevented him spending much time on meteorology. During that year, the team members began to drift apart as they realised his driving force had gone.

Charney, one of his protégées, said in 1955:

> 'The advent of the large-scale electronic computer has given a profound stimulus to the science of meteorology. For the first time the meteorologist possesses a mathematical apparatus capable of dealing with the large number of parameters required for determining the state of the atmosphere and of solving the nonlinear equations governing its motion... The radical alteration is ... [due] even more to its ability to serve as an inductive device... The machine, by reducing the mathematical difficulties involved in carrying a physical argument to its logical conclusion, makes possible the making and testing of physical hypotheses in a field where controlled experiment is still visionary and model experiment difficult, and so permits a wider use of inductive methods.'

The intersection of new computers and weather forecasting affected analysis of other complex nonlinear system as scientists and engineers drew parallels between different systems from turbulence to astrophysics to the nervous system (which Wiener described as 'a sort of cerebral meteorology').

One of the issues von Neumann grappled with was 'error amplification' and 'round-off errors'.

> 'Error amplification is ... most critical and decisive in the stepwise integration of partial differential equations (and their systems)... [However] the question of the amplification of errors is one which receives little or no consideration in the most familiar computing practices...'

In 1961, Lorenz, who had briefly studied under von Neumann (who had died in 1957), was playing with one of the descendants of von Neumann's project, a computerised weather model. Lorenz described (1993) his moment of serendipity in 1961.

---

[279] In 1942, he attended a Naval meeting on Operational Research methods for underwater mine warfare in Chicago and there met Rossby, a leading meteorologist, who talked to him about weather forecasting and thereafter continued to keep him in touch with developments in the field (Aspray p.130).

'At one point I decided to repeat some of the computations in order to examine what was happening in greater detail. I stopped the computer, typed in a line of numbers that it had printed out a while earlier, and set it running again. I went down the hall for a cup of coffee and returned after about an hour, during which time the computer had simulated about two months of weather. The numbers being printed were nothing like the old ones. I immediately suspected [some kind of "computer trouble"] but before calling for service I decided to see just where the mistake had occurred… Instead of a sudden break, I found that the new values at first repeated the old ones, but soon afterward differed by one and then several units in the last decimal place, and then began to differ in the next to last place and then in the place before that. In fact, the differences more or less steadily doubled in size every four days or so, until all resemblance with the original output disappeared somewhere in the second month. This was enough to tell me what had happened: the numbers that I had typed in were not the exact original numbers, but were the rounded-off values that had appeared in the original printout. The original round-off errors were the culprits; they were steadily amplifying until they dominated the solution. In today's terminology there was chaos.

'It soon struck me that, if the real atmosphere behaved like the simple model, long-range forecasting would be impossible. [The variables of real weather] are certainly not measured accurately to three decimal places, and, even if they could be, the interpolations between observing sites would not have similar accuracy…'

The hope had been that one could dispense with 'trivial' effects for weather systems in the same way one could for the orbits of planets. Lorenz realised that this was wrong.

Lorenz's *deterministic* weather equations programmed into a computer produced a mathematical structure now called 'chaotic' that was (i) *sensitively dependent on initial conditions*, (ii) *nonlinear*, (iii) *aperiodic*, and (iv) when turned into a graphical display showed something that came to be called a *Lorenz Attractor* (an 'attractor' is a set of states towards which neighbouring states are 'attracted' in the course of dynamic evolution[280] and a Lorenz Attractor happens also to look like a butterfly, thus invoking the famous 'butterfly effect' metaphor). In this system, cycles do not repeat exactly and *a tiny change in the input could produce a massive change in the output*. (*All* chaotic systems are nonlinear; some nonlinear systems are *not* chaotic.)

'A dynamical system is said to be chaotic if it exponentially amplifies a small difference in its conditions, so that its macroscopic behaviour appears random, even though the underlying equation is entirely deterministic' (Bennett).

Lorenz published his findings as '*Deterministic Nonperiodic Flow*' *(1963)*. He did not use the term 'chaos' and his work was mostly ignored outside a small circle of meteorologists for at least a decade. Lorenz first used the phrase 'the butterfly effect' (made particularly famous by *Jurassic Park*) in 1972 in a talk titled '*Does the Flap of a Butterfly's Wings in Brazil Set Off a Tornado in Texas?*'. *(*Ray Bradbury wrote a short story before this in which the failure of a butterfly to reproduce changed the outcome of a presidential election.) In 1972, Yorke, a mathematician, was given a copy of Lorenz's 1963 paper. He sent it to Smale who sent copies to various people studying dynamical systems. The term 'chaos' came into widespread use by scientists after a paper ('*Period Three Implies Chaos*') by Li and Yorke in 1975, which proved that certain types of deterministic equations inevitably produce chaotic results.[281] In 1976, Robert May wrote a piece for *Nature* ('*Simple*

---

[280] Ruelle and Takens' paper "*On the Nature of Turbulence*" (1971) first used the phrase "strange attractor". A strange attractor has a fractal dimension. In 2002 it was proved that a Lorenz Attractor is a Strange Attractor.

[281] It transpired that Yorke's result was a specific example of a general theory worked out a decade earlier by the Soviet mathematician Sarkovski in *Coexistence of Cycles of a Continuous Map*, 1964.

Mathematical Models With Very Complicated Dynamics') discussing *deterministic chaos*. Mandelbrot's *'The Fractal Geometry of Nature'* was published in 1977. It was roughly in the mid-1970's that 'chaos theory' took off among scientists in many different fields. It seeped into the wider culture in the 1980's after Gleick's *Chaos: The Amazing Science of the Unpredictable* (1987) became a best-seller. Lorenz wrote (1993): 'I have often speculated as to how well James Gleick's best-seller would have fared at the bookstores if it had borne a title like *Sensitive Dependence: Making a New Science*.'

Robert May's 1976 essay called on scientists to make new efforts to study nonlinear systems:

> 'Recent studies have, however, shown that the very simplest nonlinear difference equations can possess an extraordinarily rich spectrum of dynamical behaviour, from stable points, through cascades of stable cycles, to a regime in which the behaviour (although fully deterministic) is in many respects 'chaotic', or indistinguishable from the sample function of a random process…

> 'The fact that the simple and deterministic equation … can possess dynamical trajectories which look like some sort of random noise has disturbing practical implications…

> 'Alternatively, it may be observed that in the chaotic regime arbitrarily close initial conditions can lead to trajectories which, after a sufficiently long time, diverge widely. This means that, even if we have a simple model in which all the parameters are determined exactly, long term prediction is nevertheless impossible. In a meteorological context, Lorenz has called this general phenomenon the 'butterfly effect': even if the atmosphere could be described by a deterministic model in which all parameters were known, the fluttering of a butterfly's wings could alter the initial conditions, and thus (in the chaotic regime) alter the long term prediction…

> 'The elegant body of mathematical theory pertaining to linear systems (Fourier analysis, orthogonal functions, and so on), and its successful application to many fundamentally linear problems in the physical sciences, tends to dominate even moderately advanced University courses in mathematics and theoretical physics. The mathematical intuition so developed ill equips the student to confront the bizarre behaviour exhibited by the simplest of discrete nonlinear systems… Yet such nonlinear systems are surely the rule, not the exception, outside the physical sciences…

> 'Not only in research, but also in the everyday world of politics and economics, we would all be better off if more people realized that simple nonlinear systems do not necessarily possess simple dynamical properties.'

May also drew attention to the *Logistic Map*, a simple-looking equation that is often used to explain chaos.[282] For example, if we consider a population of rabbits that grows according to a simple formula, such as doubling every year, then it will keep growing in a predictable way and if we divide the population into two and watch the behaviour of each part on two different islands, we will see that we can just add the two parts back together again to get a full picture: *the whole is the sum of its parts*.

However, if the population varies according to a certain type of formula, known as a *logistic model*, then the two divided populations evolve differently and one cannot add them together to get a full picture: *the whole is more than the sum of its parts.*

---

[282] This explanation of the Logistic Map comes mainly from Mitchell (2009).

[The Logistic Map equation](#) is:
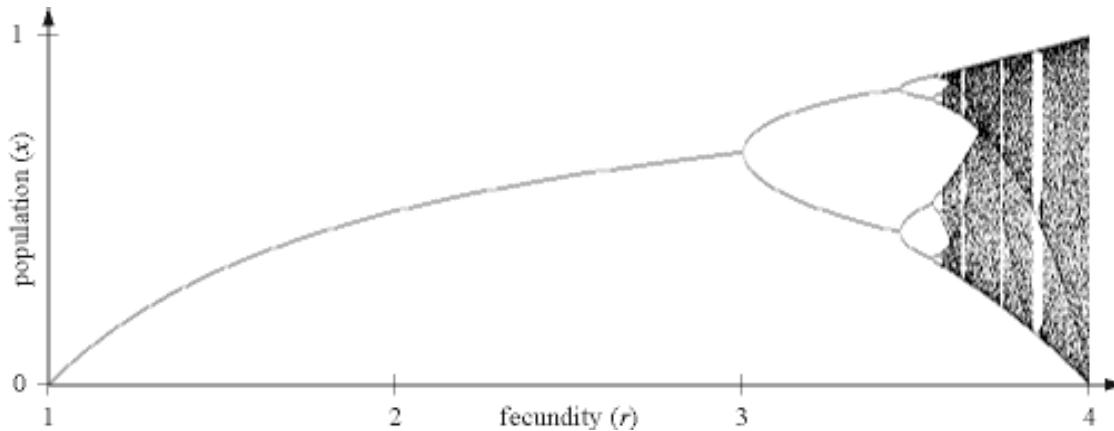
$$x_{t+1} = r * x_t (1-x_t)$$
where x is a number between 0-1 and r is a positive number

It had been studied by von Neumann, Ulam and others but it was May's article that brought it to wide attention.

As R increases from zero to about R=3.569946, the system is roughly predictable in that it moves to either a *fixed point* or a *periodic oscillation* ('fixed point' and 'oscillation' are types of *attractor*, the latter meaning that the population oscillates between two or more fixed values). At this point, x no longer settles into an oscillation around an attractor; it becomes *chaotic* (a *chaotic attractor* is sometimes called a *strange attractor*). If the input is varied by 0.0000000001, after 30 iterations two trajectories diverge significantly and soon after there is no correlation; if there is any uncertainty at all in the initial input, there is a time when the future value is unpredictable. So a simple system controlled by a deterministic equation, in which every x_t maps to only one value of x_{t+1}, can yield chaos sufficiently random that it is used to generate random numbers.

On one hand, therefore, chaotic systems are in principle beyond 'long-term' (whatever this means in the particular example) prediction. However, various properties of chaotic systems have been discovered. Some systems, including the logistic map, proceed to chaos on a predictable route - 'the period-doubling route to chaos'. Feigenbaum discovered that the rate at which bifurcations occur in the logistic map is [a constant](#) (c. 4.6692016). He then realised that all sorts of other systems (including the sine map) displayed the same constant, and gave a mathematical explanation of the various 'period-doubling cascades' that have been observed in many systems. The discovery of chaos in a system does not, therefore, mean that all hope must be abandoned of discovering useful patterns in the system.

Diagram: Bifurcation Diagram of the Logistic Map...

Endnote: Models, power laws, prediction.

'When we weave to build, / We first survey the Plot, then draw the Modell' (Shakespeare).

'... the sciences do not try to explain, they hardly even try to interpret, they mainly make models. By a model is meant a mathematical construct which, with the addition of certain verbal interpretations,

describes observed phenomena. The justification of such a mathematical construct is solely and precisely that it is expected to work – that is, correctly to describe phenomena from a reasonably wide area. Furthermore, it must satisfy certain aesthetic criteria… it must be rather simple' (von Neumann, Method in the Physical Sciences, 1955).

There are 1) *insight* models (giving experts insight into a hard problem), such as a cartoon drawn on a napkin or a Red Team wargame, and 2) *prediction* models, with the latter dividing into a) probabilistic risk assessments and b) specific point predictions.

Various (like Taleb) are obviously right to warn, echoing Mandelbrot decades earlier, about the dangers of using models that wrongly assume a 'normal distribution' of events. However, Taleb pushes his argument too far. Eg. He claims that 9/11 was fundamentally unforeseeable but terrorist events follow a 'power law' (like the distribution of nodes on the internet) and the data shows that the probability of an event of roughly 9/11 magnitude or greater between 1968-2006 was ~23%. (Extrapolating, the likelihood of another 9/11 event in the next decade is ~7%.)

We can model the frequency/magnitude of many events and examine correlations. Many things follow a 'power law'. E.g. We cannot predict individual physical accidents but we can see empirical relationships between the frequency/magnitude of smaller to larger events, and we apply the idea that eliminating smaller accidents will reduce catastrophic accidents. We cannot predict individual earthquakes but we can see empirical relationships between the frequency/magnitude of smaller to larger events. We cannot predict individual Near Earth Objects perfectly, but we can can see empirical relationships between the frequency/magnitude of smaller to larger events.

There are some obvious caveats. 1) When we do not understand the underlying physics, we are only looking at a probabilistic model and we cannot be sure why we see the correlations and when they may break down. Power law models based on a large number of small to medium scale events will not necessarily give useful information about rare events; the smaller events may not share causal similarity with the rare event. 2) In human conflicts, one has an adaptive opponent, unlike when dealing with physical accidents, earthquakes or NEOs. The collaborative approaches that are so useful with earthquakes or NEOs are much harder to organise. Red Teams are 'basically the only viable approach for understanding these feedbacks [i.e. adaptive responses].'

A problem with collaborative games is that they are usually constructed on the assumption that all the important information is known by the players and their job is to find a way of collaborating so that the information is shared and therefore the problem is solved, but this may not be the case.

> 'Much of the work on rare terrorist events seems to take for granted that "the truth is out there" and we can discover it in a sufficiently timely fashion with the right mixture of motivational assessment, social network analysis, capability measures etc. Perhaps this is the only approach which is politically defendable, but it may be the case that achieving a high probability of detection in the presence of an immense background and in the presence of informed terrorist attempts at signal concealment is simply not attainable without having such a high false positive rate as to make normal life impossible. The lack of attention to any quantitative measures of this problem makes it impossible at present to assess the likelihood of this pessimistic possibility.'

'Rare Events', JASON (2009).

An example of an 'insight model' is a recent study for the NRO (*Proteus: Insights from 2020*, 2006, updated periodically, partially classified) which predicted future conflict conducted in five

dimensions: (1) Earth, (2) Space, (3) the electromagnetic spectrum, (4) the Virtual world, and (5) Psyche.

Endnote: Constant failure

Everybody is told to 'pursue excellence', be open to criticism, don't isolate yourself from reality etc. However, almost nobody can actually do it and escape the usual psychological and bureaucratic barriers to organising high performance.

Many studies have shown that most companies fail in a few decades: 10 of the 1912 'top 100' had gone in a decade; over half vanished by 1995 (Leslie Hannah). 10% of US companies close every year. A McKinsey study of companies that had put 'quality-improvement programs' in place found that two-thirds abandoned them as failures. Just two years after Tom Peters' business bestseller '*In Search of Excellence*', a third of those identified as 'excellent' were in financial trouble. Accurate prediction of future trends is no protection: 'In the past, it required no brilliance for people to foresee the fabulous growth that awaited such industries as autos (in 1910), aircraft (in 1930) and television sets (in 1950). But the future then also included competitive dynamics that would decimate almost all of the companies entering those industries. Even the survivors tended to come away bleeding' (Buffett). Not only did Gutenberg himself go broke but so did most of the first printing companies. Only a tiny number of the thousands of early auto companies survived. Xerox's research defined much of the modern PC era but Xerox failed itself. IBM briefly succeeded with the PC but failed to understand Microsoft's challenge. Microsoft triumphed over Apple in the 1980s but was itself surprised by the internet and then by Google, which in turn was surprised by Facebook.

Large company data sets show that the extinction pattern for companies is similar to that for biological species. Ormerod argues that a 'learning model' (i.e. assuming companies learn and adapt) is incompatible with the actual signature of real extinction data, suggesting that since the real company data matches biological evolution, which proceeds with random variation (mutation and reproduction), companies do not generally display successful planning (Harford). [283]

Some areas of human activity have seen more success than others. The annual avoidable failure rate for commercial buildings in America ('building failure') is less than 0.00002%, and the airline industry built checklists that kept fatal disasters very low. However, medical research shows that a large fraction of deaths and major complications are avoidable. Medicine consistently struggled to learn how to avoid failure though the spread of evidence-based medicine over the past twenty years has brought big changes. . Although some initiatives have had great success (e.g Gawande's medical 'checklists' based on the airline industry), it is often hard to define 'best practice' properly: e.g. Medicare specified it was 'best practice' to control tightly blood sugar levels in critically ill patients but it turned out to be dangerous. 'In large part, the panels made a conceptual error. They did not distinguish between medical practices that can be standardized and not significantly altered by the condition of the individual patient, and those that must be adapted to a particular person. For instance, inserting an intravenous catheter into a blood vessel involves essentially the same set of procedures for everyone in order to assure that the catheter does not cause infection. Here is an example of how studies of comparative effectiveness can readily prove the value of an approach by which "one size fits all"… But once we depart from such mechanical procedures and impose a single "best practice" on a complex malady, our treatment is too often inadequate' (Groopman, NY Review of Books, March 2010). Most of those in politics want everything to be as relatively simple as manufacturing Coke, which involves uniform high quality and practically no failure, but in more

---

[283] This recent paper suggests why firms persist with a failing status quo.

complex systems *variation* is inevitable and good, though this also means failure is inevitable (and therefore needs to be *survivable*).[284]

Endnote: Networks, modularity, and fractals

Scaling, fractals, and metabolic networks: does fractal geometry give life an extra dimension…?

An organism's *metabolic rate* is the rate at which its cells convert nutrients to energy, this process produces heat, therefore the rate can be inferred from measuring heat. How does the *basal metabolic rate* (MR) scale with an organism's mass?

It has been known since the 1880's that smaller animals have a faster MR. If the MR scaled linearly with mass, then a hamster with eight times the mass of a mouse would have eight times the MR, and a hippo with 125,000 times the mass of a mouse would have a MR 125,000 times higher. This would mean eight or 125,000 times as much heat.

However, as the *volume* of an animal grows, the *surface area* grows slower (as Galileo worked out). The volume of a sphere = $4/3\pi r^3$. The surface area of a sphere = $4\pi r^2$. Volume scales (i.e. 'is proportional to') as the cube of the radius but surface area scales as the square of the radius; thus the surface area scales as the volume raised to the ⅔ power (i.e. $x^{2/3} = \sqrt[3]{x^2}$). If a hamster has twice the radius, it has four times the surface area and eight times the volume; if a hippo has 50 times the radius, it has 2,500 times the surface area and 125,000 times the volume.

Therefore, if the metabolic rate scaled linearly with mass and volume, then animals would have to disperse the amount of *heat cubed* but through a surface area only *squared*, and this would soon become impossible. Nature has therefore evolved a different solution. Rubner suggested in the 1880's that MR should scale with body mass in the same way as surface area, i.e. that MR scales with mass to the ⅔ power. This was generally accepted for fifty years.

However, in the 1930's, Kleiber realised that MR actually scales to body mass by the ¾ power: MR is proportional to bodymass$^{3/4}$. Widespread studies have confirmed that animals do indeed follow this power law scaling. We have also found that other biological features have a power law scaling including size:lifespan, size:heart size.

In the 1990's, two biologists, Brown and Enquist, and a physicist, Geoffrey West (previously from Los Alamos), investigated this scaling problem (henceforth the three are 'BEW').[285] They examined two questions:

1. Why does metabolic scaling follow a power law?.
2. Why does it follow the power law with exponent ¾ in particular?

Their hypothesis was that the power law was based on the fact that the metabolic system is based on *branching structures* since cells are basically the same size in hippos as in mice but hippos have more of them.

---

[284] Often it is assumed that outsiders know 'best practice' and can arrive and impose it on locals or bribe them to adopt it. Leaving aside that the 'best practice' may actually be wrong, even where it is right such attempts often misfire as locals reject outside interference. Instead, programmes can thrive by examining what a few local people have managed to do in difficult circumstances and trying to spread *their* practice instead of importing new ideas. Cf. NYT.

[285] They also came together via the Santa Fe Institute. '*The Fourth Dimension of Life*', Science, Vol. 284 no. 5420 pp. 1677-1679.

In the Koch curve, at each level the line segments are one-third the length of the previous level and the structure at each level consists of four copies of the structure at the previous level. This means the Koch curve has a fractal dimension of $3^{dimension} = 4$, therefore its dimension = c. 1.26 (i.e. log 4 / log 3). Generally, if each level is scaled by a factor of x from the previous level and is made up of N copies of the previous level, then $x^{dimension} = N$ (therefore the dimension = log N / log x).

> 'Power law distributions … *are* fractals - they are self-similar at all scales of magnification, and a power law's exponent gives the dimension of the corresponding fractal, where the dimension quantifies precisely how the distribution's self-similarity scales with level of magnification… [F]ractal structure is one way to generate a power-law distribution; and if you happen to see some quantity … follows a power-law distribution, then you can hypothesize that there is something about the underlying system that is self-similar or fractal-like.' (Mitchell, p. 265)

BEW thought that the metabolic rate was determined by the fuel efficiency of how fuel is delivered to cells therefore they examined its network structure. They assumed that circulatory and other transport networks evolve to be 'maximally space filling', that is, they can deliver efficiently to all parts of the body, and that the 'terminal units', i.e. the cells, do *not* scale with mass (which is an observed fact). Geometry shows that fractals are the best for maximally space filling objects. BEW discovered that the rate at which fuel is delivered to cells scales with body mass to the power ¾.

Rubner had suggested an exponent of ⅔ based on the scaling of mass:surface area:volume.

> 'One way to look at the ¾ exponent is that it would be result of the surface hypothesis applied to four dimensional creatures… A two-dimensional object such as a circle has a circumference and an area. In three dimensions, these correspond to surface area and volume. In four dimensions, surface area and volume correspond, respectively, to 'surface' volume and what we might call *hypervolume* - a quantity that is hard to imagine… Using arguments that are analogous to the discussion of how surface area scales with volume to the ⅔ power… [BEW are saying that] evolution structured our circulatory systems as fractal networks to approximate a 'fourth dimension' so as to make our metabolisms more efficient.'

BEW concluded:

> 'Although living things occupy a three-dimensional space, their internal physiology and anatomy operate as if they were four dimensional… Fractal geometry has literally given life an added dimension.'

BEW think their theory applies to many other biological phenomena including the fractal-like vascular networks in plants, the rate of change of DNA in organisms, and the observed ¾ scaling of species population density with body size in some ecosystems. Many others have applied the theory to other phenomena. Given the centrality of metabolism, it is naturally regarded as *a potential unifying theory in biology*.[286]

There have inevitably been challenges. Some argue that ¼ scaling laws are not as universal as BEW have claimed. Some argue that the messiness of data suggest that BEW's theory is unlike Newton's laws of motion which delivered precise predictions. Much of the arguments concern advanced statistics.

---

[286] 'I can't tell you how satisfying this was. Sometimes, I look out at nature and I think, Everything here is obeying my conjecture. It's a wonderfully narcissistic feeling.' West, NYT, 17/12/2010.

'There are always going to be people who say, 'What about the crayfish?' Well, what about it? Every fundamental law has exceptions. But you still need the law or else all you have is observations that don't make sense. And that's not science. That's just taking notes.' West

In 2002, West turned his attention to cities.[287] As cities double in size, their resource requirements grow by only about 85% yet every measure of economic activity *increases* by about 15 percent per capita but so do bad things like crime. While in biology animals slow down as they get bigger, because of increased energy demands, cities allow everything produced by human social networks to grow faster - the equivalent to finding an elephant proportionately faster than a mouse. In advanced societies, we now use about 11,000 watts per day (compared to c. 250 for a hunter-gatherer) which is more than a blue whale. We have kept innovating at an ever accelerating rate to feed this cycle but it is not sustainable, argues West.

'The scientists downloaded huge files from the Census Bureau, learned about the intricacies of German infrastructure and bought a thick and expensive almanac featuring the provincial cities of China... They looked at a dizzying array of variables, from the total amount of electrical wire in Frankfurt to the number of college graduates in Boise. They amassed stats on gas stations and personal income, flu outbreaks and homicides, coffee shops and the walking speed of pedestrians...

'After two years of analysis, West and Bettencourt discovered that all of these urban variables could be described by a few exquisitely simple equations. For example, if they know the population of a metropolitan area in a given country, they can estimate, with approximately 85 percent accuracy, its average income and the dimensions of its sewer system. These are the laws, they say, that automatically emerge whenever people "agglomerate," cramming themselves into apartment buildings and subway cars. It doesn't matter if the place is Manhattan or Manhattan, Kan.: the urban patterns remain the same. West isn't shy about describing the magnitude of this accomplishment. "What we found are the constants that describe every city," he says. "I can take these laws and make precise predictions about the number of violent crimes and the surface area of roads in a city in Japan with 200,000 people. I don't know anything about this city or even where it is or its history, but I can tell you all about it. And the reason I can do that is because every city is really the same." After a pause, as if reflecting on his hyperbole, West adds: "Look, we all know that every city is unique. That's all we talk about when we talk about cities, those things that make New York different from L.A., or Tokyo different from Albuquerque. But focusing on those differences misses the point. Sure, there are differences, but different from what? We've found the what."'

West considered the city in the same way as he had metabolism: as an energy network. Cities create economies of scale like larger animals. Infrastructure and consumption data, such as the number of gas stations or the total surface area of roads, 'showed that when a city doubles in size, it requires an increase in resources of only 85 percent.'

Initially, West says he was captivated by the similarity with his biological discoveries and was 'quite stupid'. In fact, as he soon realised, although this discovery helps suggest reasons for the evolution and spread of cities, this is not the most important thing about them. The most important thing is how they help humans' social networks.

---

[287] This section comes from a NYT piece, 12/2010.

'According to the data, whenever a city doubles in size, every measure of economic activity, from construction spending to the amount of bank deposits, *increases* by approximately 15 percent per capita. It doesn't matter how big the city is; the law remains the same. "This remarkable equation is why people move to the big city," West says. "Because you can take the same person, and if you just move them to a city that's twice as big, then all of a sudden they'll do 15 percent more of everything that we can measure." While Jacobs could only speculate on the value of our urban interactions, West insists that he has found a way to "scientifically confirm" her conjectures. "One of my favorite compliments is when people come up to me and say, 'You have done what Jane Jacobs would have done, if only she could do mathematics,' " West says. "What the data clearly shows, and what she was clever enough to anticipate, is that when people come together, they become much more productive."

'… When Bettencourt and West analyzed the negative variables of urban life, like crime and disease, they discovered that the exact same mathematical equation applied. After a city doubles in size, it also experiences a 15 percent per capita increase in violent crimes, traffic and AIDS cases. (Of course, these trends are only true in general. Some cities can bend the equations with additional cops or strict pollution regulations.) "What this tells you is that you can't get the economic growth without a parallel growth in the spread of things we don't want," Bettencourt says. "When you double the population, everything that's related to the social network goes up by the same percentage."

'… West and Bettencourt refer to this phenomenon as "superlinear scaling," which is a fancy way of describing the increased output of people living in big cities… According to West, these superlinear patterns demonstrate why cities are one of the single most important inventions in human history. They are the idea, he says, that enabled our economic potential and unleashed our ingenuity. "When we started living in cities, we did something that had never happened before in the history of life," West says. "We broke away from the equations of biology, all of which are sublinear. Every other creature gets slower as it gets bigger. That's why the elephant plods along. But in cities, the opposite happens. As cities get bigger, everything starts accelerating. There is no equivalent for this in nature. It would be like finding an elephant that's proportionally faster than a mouse."

Bigger animals require more energy therefore they cannot run around like rats. But the superlinear growth of cities has no similar constraints and instead the growth of cities feeds ever more resource consumption that more than outweighs the gains of energy efficiency.

'A human being at rest runs on 90 watts. That's how much power you need just to lie down. And if you're a hunter-gatherer and you live in the Amazon, you'll need about 250 watts. That's how much energy it takes to run about and find food. So how much energy does our lifestyle [in America] require? Well, when you add up all our calories and then you add up the energy needed to run the computer and the air-conditioner, you get an incredibly large number, somewhere around 11,000 watts. Now you can ask yourself: What kind of animal requires 11,000 watts to live? And what you find is that we have created a lifestyle where we need more watts than a blue whale. We require more energy than the biggest animal that has ever existed. That is why our lifestyle is unsustainable. We can't have seven billion blue whales on this planet. It's not even clear that we can afford to have 300 million blue whales… The only thing that stops the superlinear equations is when we run out of something we need. And so the growth slows down. If nothing else changes, the system will eventually start to collapse.'" (West)

How do we escape? Innovation. However, every resource is consumed faster so we need a faster cycle of innovations too.

> 'It's like we're on the edge of a cliff, about to run out of something, and then we find a new way of creating wealth. That means we can start to climb again… It's like being on a treadmill that keeps on getting faster. We used to get a big revolution every few thousand years. And then it took us a century to go from the steam engine to the internal-combustion engine. Now we're down to about 15 years between big innovations. What this means is that, for the first time ever, people are living through multiple revolutions. And this all comes from cities. Once we started to urbanize, we put ourselves on this treadmill. We traded away stability for growth. And growth requires change.' (West)

West has also looked at the issue of why companies' lifespans are usually so fleeting (unlike cities').

> 'After buying data on more than 23,000 publicly traded companies, Bettencourt and West discovered that corporate productivity, unlike urban productivity, was entirely sublinear. As the number of employees grows, the amount of profit per employee shrinks. West gets giddy when he shows me the linear regression charts. "Look at this bloody plot," he says. "It's ridiculous how well the points line up." The graph reflects the bleak reality of corporate growth, in which efficiencies of scale are almost always outweighed by the burdens of bureaucracy. "When a company starts out, it's all about the new idea," West says. "And then, if the company gets lucky, the idea takes off. Everybody is happy and rich. But then management starts worrying about the bottom line, and so all these people are hired to keep track of the paper clips. This is the beginning of the end."

> 'The danger, West says, is that the inevitable decline in profit per employee makes large companies increasingly vulnerable to market volatility. Since the company now has to support an expensive staff — overhead costs increase with size — even a minor disturbance can lead to significant losses. As West puts it, "Companies are killed by their need to keep on getting bigger."

> 'Unlike companies, which are managed in a top-down fashion by a team of highly paid executives, cities are unruly places, largely immune to the desires of politicians and planners. "Think about how powerless a mayor is," West says. "They can't tell people where to live or what to do or who to talk to. Cities can't be managed, and that's what keeps them so vibrant. They're just these insane masses of people, bumping into each other and maybe sharing an idea or two. It's the freedom of the city that keeps it alive."'

A 2013 paper (*The scaling of human interactions with city size*) by West *et al* shows how communication networks and intensity grow superlinearly with city population size, providing a partial explanation for the superlinear growth in economic activity. Also cf. *The Origins of Scaling in Cities* by Bettencourt, *Science* (2013).

Evolved modularity

An example of recent learning about network dynamics and the performance of complex nonlinear networks concerns *evolved modularity*. Modularity in networks improves evolvability but we have not pinned down why modularity evolves. A recent experiment applied an evolutionary algorithm to networks and found that a selection pressure to minimise the network's connection costs leads to evolved modularity.

'A long-standing open question in biology is how populations are capable of rapidly adapting to novel environments, a trait called evolvability. A major contributor to evolvability is the fact that many biological entities are modular, especially the many biological processes and structures that can be modeled as networks, such as brains, metabolic pathways, gene regulation and protein interactions. Networks are modular if they contain highly connected clusters of nodes that are sparsely connected to nodes in other clusters. Despite its importance and decades of research, there is no agreement on why modularity evolves. Intuitively, modular systems seem more adaptable, a lesson well-known to human engineers, because it is easier to rewire a modular network with functional subunits than an entangled, monolithic network...

'While most hypotheses assume indirect selection for evolvability, here we demonstrate that the ubiquitous, direct selection pressure to reduce the cost of connections between network nodes causes the emergence of modular networks. Experiments with selection pressures to maximize network performance and minimize connection costs yield networks that are significantly more modular and more evolvable than control experiments that only select for performance...

'[B]ecause the resultant modularity produces evolvability, minimizing connection costs may serve as a bootstrapping process that creates initial modularity that can then be further elevated by selection for evolvability. Such hypotheses for how modularity initially arises are needed, because selection for evolvability cannot act until enough modularity exists to increase the speed of adaptation...

'The results also open new areas of research into identifying connection costs in networks without physical connections (e.g. genetic regulatory networks) and investigating whether pressures to minimize connection costs may explain modularity in human-created networks (e.g. communication and social networks)… Knowing that selection to reduce connection costs produces modular networks will substantially advance fields that harness evolution for engineering, because a longstanding challenge therein has been evolving modular designs. It will additionally aid attempts to evolve accurate models of biological networks, which catalyze medical and biological research. The functional modularity generated also makes synthetically evolved networks easier to understand. These results will thus generate immediate benefits in many fields of applied engineering, in addition to furthering our quest to explain one of nature's predominant organizing principles.'[288]

Endnote: Some history of thinking about the limits of logic and computation, and 'P=NP?'

'But Thetis of the silver feet came unto the house of Hephaistos, imperishable, starlike, far seen among the dwellings of Immortals, a house of bronze, wrought by the crook-footed god himself. Him found she sweating in toil and busy about his bellows, for he was forging tripods twenty in all to stand around the wall of his stablished hall, and beneath the base of each he had set golden wheels, that of their own motion they might enter the assembly of the gods and again return unto his house, a marvel to look upon. Thus much were they finished that not yet were away from the fire, and gathered all his gear wherewith he worked into a silver chest; and with a sponge he wiped his face and hands and sturdy neck and shaggy

---

[288] *The evolutionary origins of modularity*, Lipson et al (2010). It also struck this author that the 'search space' in this experiment $= 10^{90}$ possible solutions (~no. of atoms in the universe) yet an evolutionary algorithm found solutions that give the right answer on 93% of the patterns while those sampled randomly from the vast search space managed a best of 23% - an example of how effective evolutionary algorithms can be.

breast, and did on his doublet, and took a stout staff and went forth limping; but there were handmaidens of gold that moved to help their lord, the semblances of living maids. In them is understanding at their hearts, in them are voice and strength, and they have skill of the immortal gods.' The Iliad, Book XVIII.

'If every tool, when ordered, or even of its own accord, could do the work that befits it … then there would be no need either of apprentices for the master workers or of slaves for the lords.' Aristotle

'The very possibility of a science of mathematics seems an insoluble contradiction. If this science is deductive only in appearance, whence does it derive that perfect rigour which no one dares to doubt? If, on the contrary, all the propositions it enunciates can be deduced one from the other by the rules of formal logic, why is not mathematics reduced to an immense tautology? The syllogism can teach us nothing that is essentially new, and, if everything is to spring from the principle of identity, everything should be capable of being reduced to it. Shall we then admit that the theorems which fill so many volumes are nothing but devious ways of saying that A is A?' Poincaré, 1896.

'Kurt Gödel's achievement in modern logic is singular and monumental – indeed it is more than a monument, it is a land mark which will remain visible far in space and time… Gödel's results showed, to the great surprise of our entire generation, that [the Hilbert Programme] could not be implemented.' Von Neumann, 1951.

'Gödel's theorem is an inexhaustible source of intellectual abuses.' Alan Sokal.

'Over the years, the constant and most reliable support of computer science - and of science generally - has been the defense establishment. While old men in congress and parliaments would debate the allocation of a few thousand dollars, farsighted generals and admirals would not hesitate to divert substantial sums to help oddballs in Princeton, Cambridge, and Los Alamos.' Metropolis, 1976.

The birth of computational thinking[289]

In Greek mythology Hephaistos gave the Cretans *Talos*, a huge mechanical statue that hurled boulders at approaching enemy ships. The Greeks possibly used steam power for mechanical devices (e.g. Archytas, a friend of Plato, reputedly built 'the Pigeon', a steam-driven model bird), Hero of Alexandria described a steam engine in the first century A.D, and we have found (but only recently understood) the amazing 'Antikythera device' - an ancient Greek analog computer used for astronomical computations.

Pascal and Leibniz…

'If I were to choose a patron saint for cybernetics[290]…, I should have to choose Leibniz. The philosophy of Leibniz centers around two closely related concepts – that of a universal symbolism and that of a calculus of reasoning. From these are descended the mathematical notation and the symbolic logic of the present day. Now, just as the calculus of arithmetic lends itself to a mechanization progressing through the abacus and the desk computing machine to the ultra-rapid computing machines of the present day, so the calculus ratiocinator *of Leibniz contains the germs of the* machina ratiocinatrix, *the reasoning machine. Indeed, Leibniz himself, like his predecessor Pascal, was interested in the construction of computing*

---

[289] Much of this section is taken from George Dyson's *Darwin Among the Machines* and *Turing's Cathedral.*

[290] In the 1940's, Wiener christened 'the entire field of control and communication theory, whether in the machine or the animal' with the name 'cybernetics which we form from the Greek χυβερνήτης or steersman'. Ampere, in '*Essay on the Philosophy of Science, or Analytic exposition of a natural classification of human knowledge*' (1843), coined the term '*cybernetique*' for the body of knowledge which originally among the Greeks referred to the steering of a ship but which they extended to 'the art of steering in general' (Dyson, p. 6).

machines in the metal. It is therefore not in the least surprising that the same intellectual impulse which has led to the development of mathematical logic has at the same time led to the ideal or actual mechanization of processes of thought.' Wiener

In 17th Century Europe, the discovery of logarithms and the invention of the slide rule helped rapid progress in calculation. John Napier developed 'logarithms' (log): if we are using base ten, then the log of 10 is 1, the log of 100 is 2, the log of 1,000 is 3 and so on; i.e. 'the log of *n*' is the power to which a base number is raised to produce *n*. The first slide rule, a mechanical calculator using logarithms, was produced by William Oughtred in 1632-3.

Pascal built a primitive calculating machine in 1642-3 and Leibniz built a computer in the 1670's but they were ignored. Leibniz was the first to set down problems of information and computation in recognisably modern form, and is often quoted as an inspiration by 20th Century pioneers in computation. First, he proposed a comprehensive and exact symbolic language that could encompass all human knowledge (*characteristica universalis*). Second, he proposed a machine (*calculus ratiocinator*) to answer questions mechanically in that language.

He described the problems humans now have in resolving disputes and understanding the world and held out a promise of what his method could bring:

> 'And although learned men have long since thought of some kind of language or universal characteristic by which all concepts and things can be put into beautiful order, and with whose help different nations might communicate their thoughts and each read in his own language what another has written in his, yet no one has attempted a language or characteristic which includes at once both the arts of *discovery* and *judgement*, that is, one whose signs and characters serve the same purpose that arithmetical signs serve for numbers, and algebraic signs for quantities taken abstractly.[291] Yet it does seem that since God has bestowed these two sciences on mankind, he has sought to notify us that a far greater secret lies hidden in our understanding, of which these are but the shadows…
>
> 'Once the characteristic numbers for most concepts have been set up, however, the human race will have a new kind of instrument which will increase the power of the mind much more than optical lenses strengthen the eyes and which will be as far superior to microscopes or telescopes as reason is superior to sight. The magnetic needle has brought no more help to sailors than this lodestar will bring to those who navigate the sea of experiments… But reason will be right beyond all doubt only when it is everywhere as clear and certain as only arithmetic has been until now.
>
> 'Then there will be an end to that burdensome raising of objections by which one person now usually plagues another and which turns so many away from the desire to reason. When one person argues, namely, his opponent, instead of examining his argument, answers generally, thus, 'How do you know that your reason is any truer than mine? What criterion of truth have you?' And if the first person persists in his argument, his hearers lack the patience to examine it. For usually many other problems have to be investigated first, and this would be the work of several weeks, following the laws of thought accepted until now. And so after much agitation, the emotions usually win out instead of reason, and we end the controversy by cutting the Gordian knot rather than untying it. This happens especially in deliberations pertaining to life, where a decision must be made; here it is given to few people to weigh the

---

[291] Leibniz sought a general problem-solving method for all questions of knowledge (*ars magna*). There were two versions of the *ars magna*: the *ars inveniendi*, which finds all true scientific statements, and the *ars iudicandi*, which allows one to decide whether it is true.

factors of expediency and inexpediency, which are often numerous on both sides, as in a balance... There is hardly anyone who could work out the entire table of pros and cons in any deliberation... And we need not be surprised that this is what has happened until now in most controversies in which the matter is not clear, that is, is not reduced to numbers.

'Now, however, our characteristic will reduce the whole to numbers, so that reasons can also be weighed.' (*On the General Characteristic*)

Elsewhere he described this new language as 'a kind of *general algebra* in which all truths of reason would be reduced to a kind of calculus.'

'... [T]his universal writing will be as easy as it is common, and will be capable of being read without any dictionary; at the same time, a fundamental knowledge of all things will be obtained. The whole of such a writing will be made of geometrical figures, as it were, and of a kind of pictures - just as the ancient Egyptians did, and the Chinese do today. Their pictures, however, are not reduced to a fixed alphabet... with the result that a tremendous strain on the memory is necessary, which is the contrary of what we propose.' (*On The Art of Combination*, 1666)

Leibniz also discovered *binary arithmetic*, which he thought would be used in his new system, and he sensed that there was something philosophically important in its properties. It was the power of binary that led him to say, '*Omnibus ex nihil ducendis sufficit unum*' (One suffices to derive all from nothing).

His most famous prediction, now often quoted at the head of chapters on computer science, was:

'What must be achieved is in fact this; that every paralogism be recognized as *an error of calculation*, and that every sophism when expressed in this new kind of notation appear as a solecism or barbarism to be corrected easily by the laws of this philosophical grammar... Once this is done, then when a controversy arises, disputation will no more be needed between two philosophers than between two computers.[292] It will suffice that, pen in hand, they sit down to their abacus and say to one another, "Let us calculate."...' 1686 (XIV)

'Yet I should venture to say that nothing more effective can well be conceived for perfecting the human mind and that if this basis for philosophizing is accepted, there will come a time, and it will be soon, when we shall have as certain knowledge of God and the mind as we now have of figures and numbers and when the invention of machines will be no more difficult than the construction of geometric problems.'

Not only did he develop the theory but he actually invented the Stepped Reckoner which used marbles and gravity to calculate in binary, apparently the world's first automated calculating machine that could perform all four arithmetical operations (the abacus is not automatic).

'Many applications will be found for this machine for it is unworthy of excellent men to lose hours like slaves in the labour of calculation which would safely be relegated to anyone else if machines were used.'

---

[292] 'Computers' here is used in the sense of human computers.

Leibniz demonstrated his 'Stepped Reckoner' to the Royal Society in 1674 but tragically it was ignored and it sat forgotten in an attic until 1879.[293]

Vaucanson's duck, punched cards, Babbage, Boole, calculating machines...

Vaucanson's duck... Jacques de Vaucanson, inspired by the Newtonian revolution, decided to build a living machine, a giant duck. It could waddle, quack, eat, and defecate and he claimed that his method had 'copied from Nature'. After it was presented to the court of Louis XV, it became internationally famous (Goethe saw it and called it 'most deplorable'). Vaucanson also created the first automated loom using punched cards in 1745. Half a century later the duck was dismantled and it was discovered that it was a con using a simple clockwork mechanism.

'The Turk'... 'The Turk' was a supposed chess automaton with the appearance of a Turk. It was made by von Kempelen and beat many celebrities at chess, including Napoleon. Von Kempelen had, like Vaucanson, conned Europe. Inside the 'automaton' was hidden a dwarf chess master.

The Jacquard Loom and punched cards... In Paris in 1801, Jacquard first demonstrated his programmable loom using punched cards to control the selection of shuttles containing coloured threads and this allowed the programming of patterns into fabrics. This was the first piece of machinery to be controlled by punched cards. Since the loom could follow instructions, the role of the skilled weaver was soon redundant. Lyon's silk weavers famously destroyed the looms but this early attempt to remove the threat of technology by forced relinquishment failed. In a decade there were thousands.

Babbage... During the Napoleonic Wars, Napoleon embarked on a project to create new maps (for property taxes) and a switch to the metric system. Both required new calculations. This required a team of low paid human calculators (some unemployed hairdressers) filling in tables line by line. It took about a decade and by then there was not the money to publish them so they languished in the *Academie des Sciences* until Charles Babbage, visiting Paris, happened to view the manuscript.

Babbage already had experience of astronomical and actuarial tables and decided to build a machine that could replace the human calculators. He proposed a Calculating Engine in 1822 and got government funding in 1824. In 1832 he produced a functioning model of the Difference Engine and published *Economy of Machinery and Manufactures*. The next year he abandoned the Difference Engine and proposed instead the Analytical Engine which would be capable of multiplying or dividing two 50-digit numbers to 100 decimal places in under a minute, and which could be programmed to evaluate polynomial expressions of any degree. It would have a processor that performed arithmetic (the mill), memory to hold numbers (the store) and the ability to alter its function *via* user input using punched cards 'adapted from those used by the card-controlled Jacquard loom' (Dyson). It was a prototype of a digital computer (a machine that applies instructions to digits to produce digits). Babbage produced thousands of pages of instructions in the hope he could get funding but the Government refused.

> 'Although framed in a dialect of gears, levers, and camshafts, Babbage anticipated the formal languages and timing diagrams that brought mechanical logic into the age of relays, vacuum tubes, transistors, microprocessors, and beyond… The analytical engine linked the

---

[293] The way in which computer scientists discuss Leibniz, the discussions of Gödel and Einstein as they strolled around Princeton discussing the views of Parmenides and Plato regarding knowledge and time, and the way in which Plato's *Forms* crop up in discussion of brain modules programmed by evolution to perform certain tasks, undermine Dawkins' view that 'pre-Darwinian answers' to questions like 'What is man?' are 'worthless' except for historical interest.

seventeenth-century visions of Hobbes and Leibniz to the twentieth century that digital computation has so transformed.' (Dyson, p.39)

Leibniz had thought that 'one could carry out the description of a machine, no matter how complicated, in characters which would be merely the letters of the alphabet, and so provide the mind with a method of knowing the machine and all its parts.' Babbage would write a century and a half later:

'By a new system of very simple signs I ultimately succeeded in rendering the most complicated machine capable of explanation almost without the aid of words… I have called this system of signs the Mechanical Notation… It has given us a new demonstrative science, namely, that of proving that any given machine can or cannot exist.'

Babbage's design anticpated Turing's demonstration that a universal computer could perform any mathematical operation. He also anticipated the idea of stored programs and discussed the library of cards for operations that could 'at any future time reproduce the calculations for which it was first arranged.' In 1991, the Science Museum built a version from his diagrams and it worked, vindicating his design.

George Boole (1815-1864)… Boole wrote '*An Investigation of the Laws of Thought, on which are founded the mathematical theories of Logic and Probabilities*' (1854). His goal was:

'… to investigate the fundamental laws of those operations of the mind by which reasoning is performed; to give expression to them in the symbolic language of a Calculus and … to make that method itself the basis of a general method for the application of the mathematical doctrine of Probabilities; and, finally, to collect from these various elements of truth … some probable intimations concerning the nature and constitution of the human mind.'

In *Boolean algebra*, the symbols +,-,x,= represent the logical operations OR, NOT, AND, IDENTITY operating on variables (x,y,z…) restricted to the binary values 0 and 1. This system brought together logic, maths, and binary. The second half of the book dealt with 'fuzzy logic' (which von Neumann would later consider in his paper '*Probabilistic Logics and the Synthesis of Reliable Organisms from Unreliable Components*') and the way in which 'individually indeterminate phenomena could nonetheless be counted on, digitally, to produce logically certain results' (Dyson). Boole also anticipated von Neumann's conclusion that the brain, using imperfect neurons, must use a statistical language to correct errors.

Punched card calculating machines… In the late 19[th] Century, others began to explore similar ideas to Babbage. In the 1880s, Allan Marquand, an art historian living in Princeton, experimented with a mechanical binary logic machine; a study of Marquand's papers suggest he may have designed the very first electrical data processor though how much of it he built is unknown. The logician Charles Peirce noticed this work and wrote a paper (*Logical Machines*, 1887) exploring 'precisely how much of the business of thinking a machine could possibly be made to perform'. The machine should be devoid of original ideas and we would no more want an original machine than 'an American board of college trustees would hire an original professor.'

The punched card industry was developed by Hollerith in the late 19[th] Century. The requirements of the 1890 US census provided an opportunity since human counting threatened to take longer than a decade. Hollerith's punched card system recorded 62 million people using 56 million cards in much more detail than had been possible. Hollerith incorporated the Tabulating Machine Company in 1896 which morphed into IBM in 1924. A 1922 *Scientific American* article discussed

how the development of punched cards 'can endow inanimate machines with brains of their own'. Dyson writes that, 'By the time Turing awakened us to their powers, the age of discrete-state machines was well under way.'

Optical data networks…

'Two distinct functions are required of a successful telegraphic code: the encoding of protocols that regulate the process of communication, and the encoding of symbols representing the message to be conveyed. Meaning … is encoded hierarchically: first by mapping elementary symbols to some kind of alphabet, then by mapping this alphabet to words, phrases, standard messages, and anything else that can be expressed by brief sequences of code. Higher levels of meaning arise as further layers of interpretation evolve. Protocols … initiate the beginning and end of a transmission and may be used to coordinate error correction and flow control.' (Dyson)

Optical data networks have a long history. For example, Clytemnestra ordered a network of beacons to bring news of the fall of Tory to her 375 miles away in Mycenae. At the beginning of Aeschylus' *Agamemnon*, the chorus asks, 'And what messenger is there that could arrive with such speed as this?' and Clytemnestra answers, 'Hephaistos, sending forth from Ida a bright radiance. And beacon ever sent beacon hither by means of the courier fire … so as to skim the back of the sea … transmitting, like a sun, its golden radiance to the look-out…' It was what would now be described as a 'one-way, one-time, and one-bit channel encoded as follows: no signal meant Troy belonged to the Trojans; a visible signal meant Troy belonged to the Greeks. Communications engineers have been improving the bandwidth ever since' (Dyson).

Polybius a thousand years later told how the Romans had improved the idea. When the Spanish Armada set sail in 1588, beacons spread the news. The telescope extended the distance between relay stations in the 17th Century. Robert Hooke wrote '*On shewing a way how to communicate one's mind at great distances*' using optical signals and a coding system that would allow signals to be exchanged securely between London and Paris in a few minutes. After scorning Leibniz's computer, put on show in 1673, Hooke built his own a few months later 'whereby in large numbers, for multiplication or division, one man may be able to do more than twenty by the common way of working arithmetic'. Unfortunately, though listed among the rarities in the collection of Gresham College in 1681, it then vanished. Hooke estimated the storage capacity of the brain by calculating the number of thoughts per second and estimated a figure of about 2 billion of which he thought about 100 million could be remembered.

In 1790, Claude Chappe tried to build an electric telegraph but abandoned it in favour of optical signals relayed by mechanical display. His prototype was destroyed twice by mobs thinking it was a scheme to signal to Louis XVI. In 1794, he built a link between Paris and Lille. In 1801, Napoleon tested a system in France for imminent use across the Channel after his invasion of England. By 1800, 'optical telegraph networks spanned most of Europe'. By 1852, the network had a total length of 3,000 miles with 3,000 operators in 556 stations about 6 miles apart. Signals could be relayed in a few seconds though in practice it ran slower. Sending a signal from Paris to Toulon took about 10 minutes. It was imitated across Europe, with the link between London and Plymouth conveying a message in three minutes.

Electric telegraph… In 1729, Stephen Gray transmitted an electric charge 765 feet. In 1747, after the invention of the Leyden jar, Watson transmitted an electric charge across the Thames. In the second half of the 18th century, there were many experiments and business ventures with electric telegraphy. In 1800, Volta announced the first electrochemical battery. In 1819, Oersted outlined the principles of electromagnetism and in 1820 Ampère gave the subject some mathematical precision.

He also discussed connecting a keyboard to an electric telegraph. In 1816, Francis Ronalds transmitted electrostatic signals over eight miles of wire and then tested his apparatus with 525 feet of insulated wire buried in underground cables. Ronalds tried to get the Navy interested but was told that 'telegraphs of any kind are now wholly unnecessary'. Many different schemes were tried and failed. In 1833, the great mathematician Gauss and Weber constructed a 1.5 mile system to share information using binary signals between the Göttingen physics department and observatory. In the1830s, various telegraph systems were deployed in Europe and America.

After a trip to Europe in 1832, Morse developed his language and the first long-distance line opened in May 1844 between Washington and Baltimore. The electric telegraph industry soon took off: in 1851, the first cable inked England and France, by 1852 there were 23,000 miles of telegraph lines, in 1861 the first line spanned North America, in 1866, the first line connected England and America, in 1870 India was connected to Europe.

> 'Telegraph signals were digital signals, whether conveyed by the on-off state of a fire beacon, the twenty-four symbol alphabet of Robert Hooke, the ninety-eight-state signal of the Chappes, a series of positive-negative voltages, or the dot-dash sequences of Morse code. To process these signals requires discrete-state machines, whether the machine is a human operator looking through a telescope and referring to page and line numbers in a book of code or one of the punched tape teleprinters that soon came to dominate telegraphy… Telegraph engineers were the first to give substance to what had been shown by Leibniz … and would be formalized by Alan Turing …: all symbols, all information, all meaning, and all intelligence that can be described in words or numbers can be encoded (and thus transmitted) as binary sequences of finite length. It makes no difference what form the symbols take; it is the number of choices between alternatives that counts. It takes five binary alternatives ($2^5 = 32$) to encode the alphabet, which is why early needle telegraphs used five separate two-state indicators and why teletypewriter tape is five holes wide. Polybius had specified two separate five-state indicators, a less efficient coding that makes sense if you are keeping a naked eye out for torches in the dark.' Dyson

The telegraph system evolved 'store-and-forward' procedures. An incoming signal to a switching node arrived as electric signal, was converted to punched paper tape identified by origin, address etc, then, after decisions by the station operators, was retranslated back to electric signal by machines that recognised the tape patterns, and sent on along the network. This was 'the ancestor of the packet-switching protocols used in computer networks today'.

> 'It was only natural that the first computers incorporated high-speed telegraphic equipment, and it is no accident that the genesis of the Colossus … was mirrored … by early steps taken toward computers by Claude Shannon... [As computer-computer communication rapidly evolved from the 1940s] no matter what the medium of exchange, the code itself and the protocols that regulate its flow remain directly descended from the first strings of telegraphic bits.' Dyson

Wireless communication… Tesla mastered wireless communication in 1893 and five years later he demonstrated the use of radio signals to control a motorboat in Madison Square Garden. Afterwards he asked the Government if they wished to buy his invention and the civil servant 'burst out laughing upon telling him what I had accomplished.' In 1917, the 'land torpedo' was built to drive explosives to enemy positions and detonate. The Kettering 'Bug', or aerial torpedo, was a small remote controlled unmanned plane using a preset gyroscope and barometer to fly to and crash into a target fifty miles away. They remained essentially prototypes and did not influence the outcome of the war. Germany used special motorboats (FL-7s) designed to be rammed into British

ships by electronic control to protect their coast. In 1916, they switched to radio waves (from cables) and this may have been the first application of Tesla's discovery to warfare (Singer).[294]

At the end of World War I, Scherbius proposed a machine that would scramble a message in one place (using a computable function) and unscramble it at the other end (by someone who knows the function) to deliver secure communications. He offered his idea to the German Navy who rejected it so he formed a company to build it himself. The 'Enigma' coding system was born and was adopted by the German Navy from 1926. (Dyson, Ch.4)

The analog 'Differential Analyser'… Until World War II, computing devices were only analog. In order to predict, say, an eclipse observatories hired teams of people ('calculators') to do the calculations by hand. One could also create an analog computer (a mechanical analog of the system) that would model the solar system with gears and shafts and run the time forward (the planetarium is also an analog computer). Before WWII the Differential Analyzer, developed by Vannevar Bush at MIT in 1929, was the leading analog computer and could approximately solve differential equations without numerical processing. 'Programming' involved setting the machine into a start state with screwdrivers, spanners and hammers.

> 'In the Differential Analyzer, motors powered shafts that turned wheels that corresponded to variables in the continuous (differential) equations defining a problem such as aiming an antiaircraft gun or operating a power grid. By selecting where the wheels rolled against other wheels, and how the shafts connected among them, the relative rates of the rotations would add up to a solution to the problem… It weighed in at a hundred tons. Programming it was an oily process, requiring days to set the ratios of the wheels and the connections of the motorized shafts… Even when everything did work, every step in [its] operation added an error of about a tenth of a percent.' (Gershenfeld)

There were about a dozen copies one of which was at the US Army's Aberdeen Proving Ground in Maryland which prepared weapons for deployment. Soldiers had to look up the angles for elevation in books of tables. Each table entry required the integration of many differential equations. A human would take days; the Differential Analyzer took about 20 minutes. In 1936, Bush predicted a shift to digital machines. The same year, Turing's famous paper was published.

The 'foundational crisis of mathematics': set theory, Hilbert, Gödel, and Turing

'What has been said once can always be repeated.' Zeno

'We must endeavour that those who are to be the principal men of our state go and learn arithmetic, not as amateurs, but they must carry on the study until they see the nature of numbers with the mind only… [A]rithmetic has a very great and elevating effect…' Plato

---
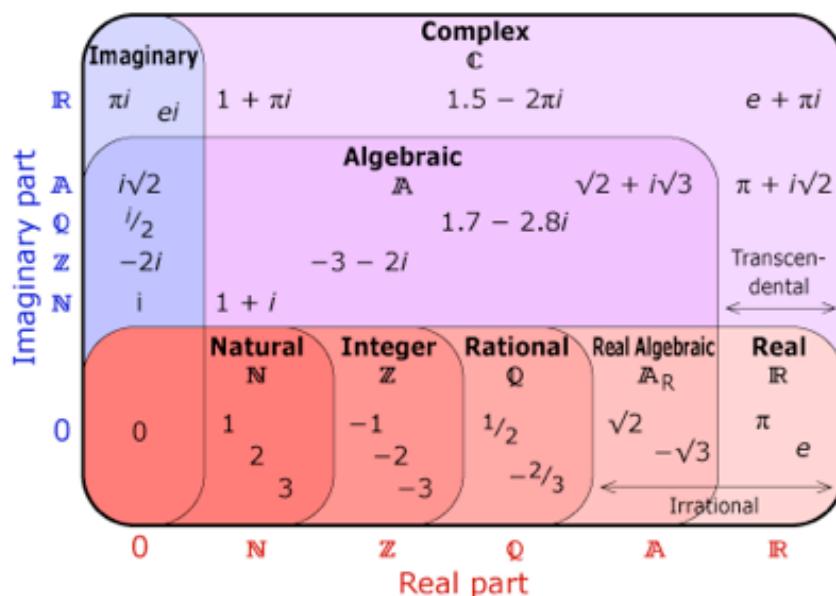
[294] In World War II, Germany fielded the first cruise missile (V1), the first ballistic missile (V2), the first jet fighter (Me-262), and the first unmanned drones - the FX-1400, a 2,300 pound bomb with four wings that would drop out of a plane and be guided to its target. (A similar American project was Operation Aphrodite, a prototype of which blew up killing Joe Kennedy's eldest son, JFK's elder brother. The Army cancelled the drone programme and JFK changed career.) The first mass produced unmanned plane was the 'Dennymite', named after a British actor who hit the big time in Hollywood and started a company after seeing a radio-controlled airplane on a movie set (Marilyn Monroe was spotted working in his factory). America's Norden bombsight was a great advance. It was an analog computer that was installed in a plane and could calculate the trajectory of bombs and drop them automatically. It was almost as expensive as the Manhattan Project and the individual sights were so valuable that they were stored in safes between missions. Norden was so cross with the US Army Air Corps that he sold it to the Navy for $1. A Norden sight was used to drop the nuclear weapons on Japan. Cf. Singer.

'I am so in favor of the actual infinite that instead of admitting that Nature abhors it, I hold that Nature makes frequent use of it everywhere.' Leibniz

'For the infinite will eventually refuse to be excluded from arithmetics…Thus we can foresee that this issue will provide for a momentous and decisive battle.' Frege

A Venn diagram of the number system

Before the 19th Century, Euclidean geometry was regarded as logically flawless with clear axiomatic foundations, while calculus was regarded as suffering an embarrassing lack of logically secure foundations. The 19th Century reversed this, as non-Euclidean geometry exploded two thousand years of assumptions and Cauchy, Weierstrass, Dedekind and Cantor created a logical foundation for calculus, though one based on taming infinity.

Around the end of the 19th Century, various schools of thought fought over the 'foundational crisis of mathematics', the problem of how to set mathematics on completely secure logical foundations. There had been 'crises' in maths before (e.g. the Pythagoreans' discovery of the irrationals, or the discovery of calculus). In the 19th Century, the discovery (invention?) of non-Euclidean geometries by Gauss and others provoked another crisis. In 1899, David Hilbert produced a new axiomatisation of Euclidean geometry (*Grundlagen der Gerometrie)* in which geometry was reduced to a purely logical treatment based on new axioms; it was logically separated from empirical concepts, or visual intuitions of space which had misled mathematicians about the logical status of Euclid's axioms. Although he used words like 'point' and 'line', the symbolic logic of the axioms was the true foundation and Hilbert would famously say that instead of words like point and line one could easily substitute 'tables, chairs, and beer mugs'. (Einstein would soon after use non-Euclidean geometry for General Relativity.)

In the 19th Century, Cantor had proved that *not all infinite sets are the same size*: the integers, rational numbers, and 'algebraic' numbers are 'denumerable'; the real numbers and transcendental numbers (like e and π) are 'non-denumerable' ('the algebraic numbers are spotted over the plane like stars against a black sky; the dense blackness is the firmament of the transcendentals' (Bell)). He developed 'set theory' to solve problems concerning infinite sets. In 1879, Frege's *Begriffsschrift*

created what is now known as 'first-order logic'.[295] Whereas Aristotle's logic could not deal with simple mathematics, Frege's purpose was to construct an axiomatic system in which all rules of inference are set out explicitly as axioms, hence results could be proved mechanically without reference to intuition.

> 'If the task of philosophy is to break the domination of words over the human mind [...], then my concept notation, being developed for these purposes, can be a useful instrument for philosophers.'

Frege was confident that his 'predicate logic' would provide the means to integrate infinity into mathematics in a consistent way and thereby end much of the recent arguments concerning the continuum and infinity. However, his confidence that his system was consistent and free from paradoxes would be famously blown up by the letter he received from Russell in 1902.

Cantor first discovered some *paradoxes* in his set theory some time in 1895 but other than a letter to Hilbert in 1896 he seems not to have discussed the issue until 1899 when he wrote to Dedekind about it. In 1897, Burali-Forti published a paper on the paradoxes of set theory but it got little attention. In 1902, just as Frege was about to finish the second volume of his *Basic Laws of Arithmetic*, he got a letter from Russell pointing out what would later become known as *Russell's paradox*: what about a *set (S) of sets that are not members of themselves*? If S is *not* a member of itself, then it is 'a set that is not a member of itself', therefore it *should* be a member of itself. (A more colloquial formulation is: a barber shaves all those who do not shave themselves (set S); is the barber a member of the set S? Either a yes or no answer gives a contradiction.)

Frege wrote, in a Postscript of October 1902 that was included in the publication of 1903:

> 'Hardly anything more unfortunate can befall a scientific writer than to have one of the foundations of his edifice shaken after the work is finished. This was the position I was placed in by a letter from Mr. Bertrand Russell just when the printing of this volume was nearing its completion…

> '*Solatium miseris, socios habuisse malorum*… What is in question is not just my particular way of establishing arithmetic but whether arithmetic can possibly be given a foundation at all.'

In response to these paradoxes, which caused such bewilderment and concern for mathematics, two schools formed: the *formalists* (whose most famous champion was Hilbert), whose response was not to abandon Cantor's set theory but to seek *a complete and consistent axiomatic system for arithmetic* to banish the paradoxes and contain problems arising from the use of infinities, and the *intuitionists* (whose most famous champion was Poincaré), who thought this attempt doomed to fail and who regarded Cantor's introduction of infinities as the main reason for the disasters.

The problem was later summarised by von Neumann (1947):

> 'In the late nineteenth and the early twentieth centuries a new branch of abstract mathematics, G Cantor's theory of sets, led into difficulties. That is, certain reasonings led to contradictions; and, while these reasonings were not in the central and 'useful' part of set theory, and always easy to spot by certain formal criteria, it was nevertheless not clear why they should be deemed less set-theoretical than the 'successful' parts of the theory. Aside from the *ex post* insight that they actually led into disaster, it was not clear what *a priori*

---

[295] Frege's First-Order Logic was based on: objects (e.g. 1); predicates, or relationships (e.g. 1<2); operations, or functions (e.g. 1+2); logical operations (e.g. AND); quantifiers (e.g. 'for all').

motivation, what consistent philosophy of the situation, would permit one to segregate them from those parts of set theory which one wanted to save. A closer study of the merits of the case, undertaken mainly by Russell and Weyl, … showed that the way in which not only set theory but also most of modern mathematics used the concepts of 'general validity' and of 'existence' was philosophically objectionable.'

Russell's *Principia* (published in three volumes 1910-13) was an attempt to rescue Frege's reduction of mathematics to logic. Russell thought that he had at least partially succeeded and even post-Gödel would write that he had cleared up 'two millennia of muddle-headedness about "existence", beginning with Plato's *Theaetetus*'. He had not. By the 1920's, Zermelo and Fraenkel had established an axiomatisation of set theory (now known as 'ZF set theory') on which to base 'normal' mathematics but a question remained - *is it complete and consistent?*

In his 1925 *On the Infinite*, Hilbert explained why his Programme was of fundamental importance:

> 'Admittedly, the present state of affairs where we run up against paradoxes is intolerable. Just think, the definitions and deductive methods which everyone learns, teaches, and uses in mathematics, lead to absurdities!...

> 'Where else would reliability and truth be found if even mathematical thinking fails?... The definitive clarification of the *nature of the infinite* has become necessary, not merely for the special interests of the individual sciences, but rather for the *honour of human understanding* itself.'

In various papers and talks between 1925-1928 he described the search for a formal system expressing 'the whole thought content of mathematics in a uniform way' and described such a system as like 'a court of arbitration, a supreme tribunal to decide fundamental questions - on a concrete basis on which everyone can agree and where every statement can be controlled'. In 1928, Hilbert made his Programme more precise by posing what became known as the *Entscheidungsproblem* – 'the decision problem'.

> 'The *Entscheidungsproblem* is solved when one knows a procedure by which one can decide in a finite number of operations whether a given logical expression is generally valid or is satisfiable. The solution of the *Entscheidungsproblem* is of fundamental importance for the theory of all fields, the theorems of which are at all capable of logical development from finitely many axioms.' (Church later wrote that the *Entscheidungsproblem* should be understood as a way 'to find an effective method by which, given any expression Q in the notation system, it can be determined whether or not Q is provable in the system.')

In a 1928 lecture in Bologna, Hilbert declared:

> 'With this new foundation of mathematics, which one can conveniently call proof theory, I believe the fundamental questions in mathematics are finally eliminated, by making every mathematical statement a concretely demonstrable and strictly derivable formula…

> '[I]n mathematics there is no *ignorabimus*, rather we are always able to answer meaningful questions; and it is established, as Aristotle perhaps anticipated, that our reason involves no mysterious arts of any kind: rather it proceeds according to formulable rules that are completely definite – and are as well the guarantee of the absolute objectivity of its judgement.'

Hilbert's approach is a reformulation of Leibniz's search for a *characteristica universalis, ars inveniendi* and *ars iudicandi*. Von Neumann described the consequence of finding such a 'decision procedure' (1927):

> '… then mathematics, in today's sense, would cease to exist; its place would be taken by a completely mechanical rule, with the aid of which any man would be able to decide, of any given statement, whether the statement can be proven or not.'

On 8 September 1930, Hilbert gave a farewell address in Königsberg and boldly stated:

> 'For the mathematician there is no *Ignoramibus*, and, in my opinion, not at all for natural science either... In an effort to give an example of an unsolvable problem, the philosopher Comte once said that science would never succeed in ascertaining the secret of the chemical composition of the bodies of the universe. A few years later this problem was solved…

> 'The true reason, according to my thinking, why Comte could not find an unsolvable problem is, in my opinion, that there is no unsolvable problem. In contrast to the foolish *Ignoramibus* ['*Ignoramus et ignoramibus*', We do not know and we shall never know], our credo avers: *Wir müssen wissen wir werden wissen* [We must know, We shall know].'[296]

Hilbert was confident that the combined attack of von Neumann, Bernays and Ackerman would very soon finally solve his Second Problem and the *Entscheidungsproblem*. He did not know that the day before an unknown Austrian, Kurt Gödel, had quietly announced his own momentous results to an initially underwhelmed conference.

Gödel attended a conference in Königsberg, '*Epistemology of the Exact Sciences*', on 5-7 September 1930. On the 6th, Gödel gave a talk on his 1930 proof (for his doctoral dissertation) of *the completeness of first order predicate calculus*.[297] On Sunday 7th September,[298] in a roundtable discussion including von Neumann and Carnap[299] on the foundations of mathematics, he first announced his First Incompleteness Theorem. None of the participants seem to have realised the importance of what he was saying with the exception von Neumann (one of the organisers did not mention the result in his report of the conference) who immediately realised the significance of what Gödel had said, spoke to him at the end of the session, and on 20th November wrote to Gödel suggesting the Second Incompleteness Theorem. Three days earlier on 17th November, Gödel's paper *On Formally Undecidable Propositions of* Principia Mathematica *and Related Systems I (1931)*,[300] had been received for publication; it was published in early 1931. Gödel replied to von

---

[296] A record of Hilbert saying this on German radio in 1930 still exists, on which one can hear him finish with "…*wir werden wissen*" and laugh as the broadcast ends. This famous phrase is engraved on his tombstone.

[297] Franzen points out that there are many misunderstandings about this because of a natural misunderstanding of two different technical meanings of 'complete'. 'That [FOPC] is complete does not mean that some formal system is complete in the sense of Gödel's incompleteness theorem [1931].'

[298] This date is often mis-stated. I have opted to trust Feferman.

[299] Comte (1798-1857), the father of Positivism, envisioned *sociology* as a science of the social that would replace metaphysical nonsense with scientific methodology and this desire caught the later 19th Century *Zeitgeist* as Europe industrialised (he had earlier used the phrase 'social physics'). After World War I, the Vienna Circle developed what they called 'Logical Positivism' out of what they thought were the lessons of Wittgenstein's *Tractatus*. One of Wittgenstein's most influential acolytes was Carnap. Von Neumann and Gödel thought that Carnap never understood the Incompleteness Theorems. Carnap was, wrote Von Neumann, a 'turbid source' of ideas; he merely expressed 'completely naïve simplistic views with a terribly important air', and Carnap 'obviously has absolutely no understanding of the meaning of Gödel's results'. Gödel would later write that Wittgenstein 'advances a completely trivial and uninteresting misinterpretation' of the Incompleteness Theorems.

[300] A Part II was planned but not written.

Neumann that he had indeed proved it. Von Neumann realised that after years of effort, Hilbert's Programme was probably doomed.

The opening paragraphs explain his revolution:[301]

> 'The development of mathematics toward greater precision has led, as is well known, to the formalization of large tracts of it, so that one can prove any theorem using nothing but a few mechanical rules. The most comprehensive formal systems that have been set up hitherto are the system of *Principia Mathematica* (PM) and the [Zermelo-Frankel] axiom system of set theory… These two systems are so comprehensive that in them all methods of proof today used in mathematics are formalized, that is, reduced to a few axioms and rules of inference. One might therefore conjecture that these axioms and rules of inference are sufficient to decide *any* mathematical question that can at all be formally expressed in these systems. It will be shown below that this is not the case, that on the contrary there are in the two systems mentioned relatively simple problems in the theory of integers that cannot be decided on the basis of the axioms. This situation is not in any way due to the special nature of the systems that have been set up but holds for a wide class of formal systems; among these … are all systems that result from [PM and ZF] through the addition of a finite number of axioms, provided no false propositions … become provable owing to the added axioms…'

Gödel later explained his results in a Postscriptum added to his 1931 paper (28 August 1963), which states the two incompleteness theorems in plain language:

> 'In consequence of later advances, in particular of the fact that due to A.M. Turing's work a precise and unquestionably adequate definition of the general notion of formal system[302] can now be given, a completely general version of Theorems VI and XI [the two incompleteness theorems of the 1931 paper] is now possible. That is, it can be proved rigorously that in *every* consistent formal system that contains a certain amount of finitary number theory there exist undecidable arithmetic propositions and that, moreover, the consistency of any such system cannot be proved in the system.'[303]

What do *consistent, complete, and decidable* mean in Gödel's analysis?

> A formal system S is *consistent* if there is no A such that both A and ~A are theorems.

> A sentence is *undecidable* if neither A nor ~A is a theorem.

> S is *complete* if no sentence in the language of S is *undecidable* in S; otherwise it is *incomplete*. (Franzen)

Feferman, a leading scholar of Gödel, summarises his 1931 results thus:

> 'If S is a formal system such that:
> (i) the language of S contains the language of arithmetic,
> (ii) S includes PA [Peano Arithmetic], and

---

301 I use the translation by van Heijenhort which is preferred by scholars partly because it was carefully checked by Gödel himself. Unfortunately another translation is more widely used leading to some problems.

302 Footnote in original: 'In my opinion the term 'formal system' or 'formalism' should never be used for anything but this notion [ie. Turing's]…' The 'characteristic property' of formal systems 'is that reasoning in them, in principle, can be completely replaced by mechanical devices.'

303 *Collected Works*, Volume I, p. 195.

(iii) S is consistent

then there is an arithmetical sentence A which is true but not provable in S [first incompleteness theorem]… [and] the consistency of S is not provable in S [second incompleteness theorem].'

Turing's *On Computable Numbers, With an Application to the Entscheidungsproblem (1936)* combined Gödel's work with the thought experiment of a particular type of 'computer', which is now known as a Turing Machine (TM), to mechanise the analysis of the *Entscheidungsproblem*.[304] Turing demonstrated the equivalence of the *Entscheidungsproblem* and the question: is there a general programme (algorithm) which can decide in advance whether any particular programme halts (terminates) or not ('the Halting Problem'). He showed that a 'Universal Turing Machine' (UTM) can simulate the action of any individual Turing Machine and it can *compute anything that can be computed using the rigorous apparatus of a suitable formal axiomatic system*. However, this does *not* mean that a UTM can solve any mathematical problem: some problems are absolutely unsolvable on any computer or any possible technology - they are *uncomputable*.

Turing's paper also rests on Cantor's work on the non-denumerability of the reals and, like Gödel, applies Cantor's diagonal method for his proof that 'computable' numbers do not include all definable numbers and are therefore, unlike the reals, denumerable and this means that a computer cannot solve the *Entscheidungsproblem*.

'The "computable" numbers may be described briefly as the real numbers whose expressions as a decimal are calculable by finite means… A number is computable if its decimal can be written down by a machine…

'I show that certain large classes of numbers are computable. They include for instance the real parts of all algebraic numbers, … the numbers π, e, etc [NB. π and e are transcendental, so the computable numbers include some, but not all, transcendentals]. The computable numbers do not, however include all definable numbers…

'Although the class of computable numbers is so great, and in many ways similar to the class of real numbers, it is nevertheless denumerable. [NB. The set of all real numbers *is* non-denumerable as Cantor proved.]…'

Von Neumann summarised Turing's conclusion for the non-mathematician:

'The problem is to determine, for a class of logical expressions or propositions, whether there is a mechanical method for deciding whether an expression of this class is true or false… Turing proved that there is something for which you cannot construct an automaton; namely, you cannot construct an automaton which can predict in how many steps another automaton, which can solve a certain problem, can actually solve it. So, you can construct an automaton which can do anything any automaton can do, but you cannot construct an automaton which will predict the behaviour of any arbitrary automaton. *In other words, you can build an organ which can do anything that can be done, but you cannot build an organ which can tell you whether it can be done… This is connected with … the results of Gödel.* The feature is just this, that you can perform within the logical type that's involved everything that's feasible,

---

[304] 'He told me that the 'main idea' of the paper came to him when he was lying in Grantchester meadows in the summer of 1935' (Gandy).

but the question of whether something is feasible within a type belongs to a higher logical type.'[305]

One of the current standard graduate textbooks on computer theory summarises: Turing proved that you cannot have a general TM that can examine any other TM and tell if it halts; it is necessarily driven into contradiction. (Sipser, p.181)

> '[You have] a computer program and a precise specification of what that program is supposed to do (e.g. sort a list of numbers). You need to verify that the program performs as specified… Because both the program and the specification are mathematically precise objects, you hope to automate the process of verification by feeding these objects into a suitably programmed computer. However, you will be disappointed. The general problem of software verification is not solvable by computer.' (Sipser)

It has been claimed by many that the Gödel-Turing results somehow prove that certain types of physical theories are impossible. This has even been claimed by prominent scientists such as Freeman Dyson and Stephen Hawking. However, the experts on Gödel-Turing say that such opinions are a misunderstanding.

For example, Freeman Dyson wrote the following:

> 'Now I claim that because of Gödel's theorem, physics is inexhaustible too. The laws of physics are a finite set of rules, and include the rules for doing mathematics, so that Gödel's theorem applies to them. The theorem implies that even within the domain of the basic equations of physics, our knowledge will always be incomplete.'[306]

Saul Feferman, chief editor of Gödel's *Collected Works*, replied pointing out that while it was true that no future physical theory could 'solve' or escape Gödel's incompleteness theorems, this is not at all the same thing as prohibiting a supposedly 'complete' set of physical laws:[307]

> '… if the laws of physics are formulated in an axiomatic system S which includes the notions and axioms of arithmetic as well as physical notions such as time, space, mass, charge, velocity, etc., and if S is consistent, then there are propositions of higher arithmetic which are undecidable by S. But this tells us nothing about the specifically physical laws encapsulated in S, which could conceivably be complete as such.'

Feferman made the further point that all the maths required by physics requires only the ZFC axiom system used for all normal maths 'and there is not the least shred of evidence' that a stronger system is needed: 'In fact, it has long been recognized that much weaker systems than that suffice for scientific applications.'

Feferman then distinguished between applied maths' untroubled use of ZFC and the metamathematical problems of pure maths and set theory.

---

[305] When the editor of his notes reviewed them after his death, he wrote to Gödel asking about the references to his work. Gödel replied: 'I have some conjecture as to what von Neumann may have had in mind… [A] complete epistemological description of a language A cannot be given in the same language A, because the concept of truth of sentences of A cannot be defined in A. It is this theorem that is the true reason for the existence of undecidable propositions in the formal systems containing arithmetic.'

[306] NY Review of Books, 13 May 2004. Feferman reply, 15 July 2004

[307] NYRB weblink above.

'It is entirely another matter whether, and in what sense, pure mathematics needs new axioms beyond those of the Zermelo-Fraenkel system; that has been a matter of some controversy among logicians.'

Dyson replied to the NYRB:

'I am grateful to Solomon Feferman for explaining why we do not need Gödel's theorem to convince us that science is inexhaustible.'

Hawking in a 2002 speech, 'Gödel and the End of Physics'[308], made a similar argument: 'Maybe it is not possible to formulate the theory of the universe in a finite number of statements. This is very reminiscent of Gödel's theorem.' Hawking argued that 'if there are mathematical results that cannot be proved, there are physical problems that cannot be predicted' and gives an example of trying to solve the Goldbach Conjecture in physical form by arranging blocks of wood.

As Franzen stresses, Gödel's and Turing's 1931/1936 only tell us about formal systems with arithmetic.

'The basic equations of physics … cannot decide every arithmetical statement [and could not and will not however they develop], but whether or not they are complete considered as a description of the physical world, and what completeness might mean in such a case, is not something that the incompleteness theorem tells us anything about…

'Our predictions of the outcome of physical experiments using arithmetic are based on the premise that arithmetic provides a good model for the behaviour of certain actual physical systems with regard to certain observable properties… The relevant description of the physical world amounts to the assumption that this premise is correct. The role of the arithmetical statements is as a premise in the application of this description to arrive at conclusions about physical systems…

'… [N]othing in the incompleteness theorem excludes the possibility of our producing a complete theory of stars, ghosts and cats, all rolled into one, as long as what we say about stars, ghosts and cats can't be interpreted as statements about the natural numbers.'[309]

It is also often claimed (most famously by Roger Penrose) that the Gödel-Turing results 'prove' that intelligent machines are impossible. Most of the leading scholars of the subject do not agree with such claims (cf. Franzen) and Turing himself rejected such arguments.

Despite the fact that 'the incompleteness theorem shows that we cannot formally specify the sum total of our mathematical knowledge' (Franzen), and proves the existence of 'undecidable propositions' and 'uncomputable numbers', it should be stressed that so far these things have been confined to the far edge of mathematical logic - no 'normal' mathematical problem has been proved to be 'undecidable'.

'… so far, no unsolved problem of prior mathematical interest like these [Goldbach Conjecture, Riemann Hypothesis] has even been shown to be independent of Peano

---

[308] http://www.damtp.cam.ac.uk/strings02/dirac/hawking/

[309] Franzen p.88-90.

Arithmetic. The true statement shown to be unprovable by Gödel is just contrived to do the job; it doesn't have mathematical interest on its own.' Feferman

'No arithmetical conjecture or problem that has occurred to mathematicians in a mathematical context, that is, outside the special field of logic and the foundations or philosophy of mathematics, has ever been proved to be undecidable in ZFC…

'… a proof that the twin prime conjecture is undecidable in ZFC would be a mathematical sensation comparable to the discovery of an advanced underground civiliazation on the planet Mars.' (Franzen)[310]

Cf. 'The Church-Turing-Deutsch principle' in the next Endnote on quantum computation.

Von Neumann and the digital computer

'The [IAS] School of Mathematics has a permanent establishment which is divided into three groups, one consisting of pure mathematics, one consisting of theoretical physicists, and one consisting of Professor von Neumann.' A young Freeman Dyson.

Atanasoff… Atanasoff was a physicist who experimented with IBM tabulators to help him solve equations faster. Struggling with a problem one day, he got in his car and drove 200 miles to the Mississippi, sat with a bourbon and soda and began sketching a solution on napkins. Between 1939 and 1940 he built working computers using vacuum tubes.

'Atanasoff's machine was a little-known computer that was restricted to a narrow class of problems, was not programmable and was never fully functional. Atanasoff discontinued development in 1942, and his work was virtually unknown until 1971, when Honeywell brought the suit against Sperry Rand to invalidate the ENIAC patent. During the trial it was revealed that Mauchly had visited Atanasoff and had seen his computer in June 1941. What he learned from this visit cannot be known, but the design of the ENIAC bore no resemblance to the Atanasoff computer. Mauchly himself claimed that he took away "no ideas whatsoever." Although the judge gave priority of invention to Atanasoff, this was a legal judgment that surprised many historians.' (Campbell-Kelly, *Scientific American*, 2009)

Bletchley… While in Princeton in 1937, where he got to know von Neumann,[311] Turing had also experimented with building primitive electric calculators and tried engineering parts himself in the Princeton labs. In England, Tommy Flowers was using vacuum tubes as electronic switches in the phone system. Both would be deeply involved in Britain's secret wartime codebreaking effort at Bletchley Park.

The first binary electro-mechanical computer was Conrad Zuse's Z3 in 1941 but it was only partly programmable. The first digital electronic and partly programmable computer was the Colossus designed by Flowers. It was tested at the end of 1943 and went live at Bletchley in early 1944. Bletchley's original 'bombes' were analog electromechanical machines containing rotors to mimic the Enigma that would tick through possibilities; silence would signal that the machine had found a clue. They were produced by the British Tabulating Machine. However, the longer messages could not be cracked using these 'bombes'. The U-boat based machines that Turing faced in Bletchley Park could take one of about 10^{23} initial states. A brute force search would have required scanning about

---

[310] In 2013 the twin prime conjecture was finally proved.

[311] Von Neumann tried to persuade Turing to work with him at the IAS but Turing decided to return to England.

200,000 states per second to complete a search within 15 billion years. Electronic punched-tape machines were built (nicknamed 'Heath Robinsons') to attack the problem.

One decoding system involved punched card machines comparing two bits of coded tape; Flowers suggested improving the process by transferring one of the sequences to a machine's *internal memory*. The new machine, Colossus, had 1,500 valves (vacuum tubes) operating with a Boolean logic program input by a plug-board and toggle switches at the back of the machine. It was not a stored-program computer (executing and modifying internally stored programs) and was not a *general purpose* ('Turing complete') computer but it was the first *partly programmable electronic digital computer* and 'it came almost as close as ENIAC and some two years in advance.' The Colossus was extremely heavily classified until the 1970s and the original documents and prototypes were destroyed. The classification of computer developments greatly undermined Britain's ability to develop a commercial computer industry. The alumni of Bletchley demonstrated a working stored-program computer, the Manchester Baby Mark I, in June 1948 but could not develop many of their ideas. (Dyson, Ch.4)

The ENIAC… After Pearl Harbor there was a huge surge in US demand for better computation.[312] In spring 1942 John Mauchly at the Moore School of Electrical Engineering (University of Pennsylvania) had the idea of making an electronic computer using vacuum tubes instead of mechanical components. After delays and snafus, the idea came to the attention of Herman Goldstine who was the liaison between Aberdeen Proving Ground and the Moore school. Goldstine rapidly got funding for the project which became known as ENIAC (Electronic Numerical Integrator and Computer) and construction began on 9 April 1943.

The life of a vacuum tube was about 3,000 hours and ENIAC's initial design required about 5,000 tubes so it would barely be able to function for a few minutes before it failed. Eckert knew that the tubes failed under the stress of being switched on/off and their life could be extended by using them below their rated voltage. They also built in redundancy by having 18,000 tubes. They finished the machine in 1945. It weighed 30 tons and consumed 200kw. It could do 5,000 additions per second and compute a trajectory in less time than it took for a shell to find its target. However, it could only store 20 numbers at a time and programming took days. *The ENIAC was the first 'Turing-complete', general purpose programmable digital electronic computer (but it was not originally a 'stored program' computer).*

In December 1945, ENIAC ran its first programme – for the hydrogen bomb project at Los Alamos. Ulam, part of the Los Alamos team, later wrote:

> 'One could hardly exaggerate the psychological importance of this work and the influence of these results on Teller himself and on … the Los Alamos laboratory in general… I well remember the spirit of exploration and of belief in the possibility of getting trustworthy

---

[312] At the start of the Manhattan Project, only about a dozen hand-operated calculating machines were available. Feynman and Metropolis fixed them to avoid long repair times and persuaded the senior people to invest more in them. The codebreakers of the US Navy's top secret OP-20-G also were limited to a handful of IBM electromechanical machines at the start of WWII. During the war, the pressure of events kept them focused on immediate problems and building a succession of analog machines to break specific codes; they did not spend resources on a general purpose programmable electronic computer. It was not until they sent James Pendergrass to the Moore School in 1946, where he learned of von Neumann's plan for the EDVAC, that OP-20-G began its own plans for an electronic computer. The Pendergrass report (1946) remained Top Secret for decades and remains only partially published.

answers in the future. This partly because of the existence of computing machines which could perform much more detailed analysis and modelling of physical problems.'[313]

That programme was written by von Neumann.

Von Neumann and the EDVAC…

In the 1920s, von Neumann had spent much of his time working with Hilbert on his Programme (above), which ceased after he heard Gödel announce his revolution at the end of 1930. In the 1930s, he spent much of his time working on the mathematical foundations of the newly developed quantum mechanics (cf. his classic '*The Mathematical Foundations of Quantum Mechanics*' and his papers on 'quantum logic'). He loathed the Nazis and soon after their rise he left for the newly created Institute of Advanced Studies (IAS) at Princeton, where Einstein and Gödel would later arrive.

Before, during and after World War II until his premature death in 1957, von Neumann worked on a combination of: a) Operations Research (OR),[314] much of which confronted him with the problems of *non-linear partial differential equations (PDEs)* in fluid dynamics, shockwaves, and turbulence (the most important were the Manhattan Project, the hydrogen bomb, and ICBMs[315]); b) the creation of Game Theory in an attempt to improve the mathematical analysis of decisions and economics; c) the creation of the modern computer (which now operates with what is known as 'von Neumann architecture') and its application to specific problems (e.g. the first computer system for early warning and response, (SAGE)); d) connections between machine intelligence and analysis of the human brain; e) 'cellular automata' and what he called a general 'logical theory of automata' and his proof (the first) of the logical possibility of self-replicating machines (which paved the way for the field of 'Artificial Life'). In autumn 1955 von Neumann was diagnosed with cancer, by early 1956 he was confined to a wheelchair, and instead of giving his planned lectures on *The Computer and the Brain* he entered Walter Reed hospital and died on 8 February 1957, after almost every senior figure in the US Government visited him to thank him for his service.

The work he did in this fifteen year period therefore linked logic and computational thinking, physics, the mind, information theory, artificial intelligence and life, economic modelling, and fundamental mathematics. All of these areas confronted him with the inadequacies of conventional mathematics for dealing with *non-linear systems* and this was one of his main motivations for the development of the computer. He hoped that computers would provide an alternative to the conventional scientific process of – observe, hypothesise, predict, experiment – by allowing

---

[313] Even such illustrious scientists as Enrico Fermi thought that ENIAC could never be built reliably enough to be useful. He was a rapid convert to the new field. Metropolis described him and Teller as 'the first hackers' who taught themselves the menial details of the repair of MANIAC so they understood it properly and could reprogram it themselves. Fermi's last paper concerned serendipitous discoveries about nonlinear systems that came from playing with MANIAC, which Metropolis called 'the first major scientific discovery made by computer and it is not fully understood to this day' (Metropolis, op.cit): *Studies of Nonlinear Problems I*, Fermi, Ulam 1954-5.

[314] OR involved the integration of scientists, engineers, the military, economists, and strategists in the study of the tactical, operational, strategic, and technological aspects of systems, especially weapons systems: e.g, the use of radar to aim anti-aircraft fire, changing the depth of depth charges, the convoy system, the correct bombing strategy of Japan etc. The RAND Corporation, the original 'think tank' established by the USAF after the war, was one of the main homes for OR. Cf. *The Wizards of Armageddon*, by Fred Kaplan, which explores the role of RAND and 'defence intellectuals' in developing nuclear strategy. Williams was influential in driving debate at RAND about preventive war: cf. *Regarding the Preventive War Doctrine*, (unpublished) RAND 1953.

[315] Von Neumann's classified *Strategic Missile Evaluation Committee* (codenamed the Teapot Committee) produced a plan adopted by Eisenhower as the military's 'highest priority'. His Committee worked on the Atlas rocket and the beginning of serious thought within the US military concerning *space*; one of the spin-offs was the first military satellite surveillance system CORONA (operating from 1959, declassified 1995).

simulations and *heuristic use of computers*, 'the methods of automatic perception, association, organization and direction'.

> 'Our present analytical methods seem unsuitable for the solution of the important problems arising in connection with non-linear partial differential equations and, in fact, with virtually all types of nonlinear problems in pure mathematics… A brief survey of … most of the successful work in pure and applied mathematics suffices to show that it deals in the main with linear problems… The advance of analysis is, at this moment, stagnant along the entire front of non-linear problems.' (*Large-Scale High Speed Computing*, vN & Goldstine)

This is not a 'transient' problem: 'we are up against an important conceptual difficulty' as demonstrated by the fact that the 'main mathematical difficulties in fluid dynamics' have been known for a long time but 'hardly any progress which could be rated as important' by reasonable criteria have been made. This failure required progress with digital computers.

> '[M]any branches of both pure and applied mathematics are in great need of computing instruments to break the present stalemate created by the failure of the purely analytical approach to the non-linear problems… [R]eally efficient high speed computing devices may, in the field of non-linear partial differential equations as well as in many other fields … provide us with those heuristic hints which are needed in all parts of mathematics for genuine progress' (ibid.)

Von Neumann first learned of ENIAC in 1944. He had travelled to England in early 1943, a trip that sparked what he described as his 'obscene interest' in computing (it seems he met Turing during this trip but the details of their conversation are likely unrecoverable). During 1943, he also read the McCulloch and Pitts paper '*A Logical Calculus of the Ideas Immanent in Nervous Activity*', one of the pioneering papers on the brain.[316] It presented the brain as a *network of neurons* and claimed to have 'proved, in substance, the equivalence of all general Turing machines – man-made or begotten'. Von Neumann wrote of it that it contained 'very good and important ideas' but was 'burdened with philosophical terminology that is not entirely clear' and the 'cumbersome and inappropriate logical apparatus of R. Carnap', of whose philosophy he strongly disapproved (as did Gödel).

His work at Los Alamos and elsewhere involved work on computation devices and in summer 1944, while crisscrossing the country, he ran into Herbert Goldstine on a train platform. He discovered that the Moore School of Electrical Engineering (University of Pennsylvania) was working on ENIAC. From this point on, von Neumann was deeply involved with the development of the ENIAC and its first program was written by him.

He also began work on a successor. In spring 1945, he wrote '*First Draft of a Report on the EDVAC*' (Electronic Discrete Variable Arithmetic Computer), one of the most influential documents in the history of computer science. He set out the logical architecture for a *stored program computer*: a *memory unit*, including for operating instructions; an *arithmetic unit* for operations; an *input unit* to put program and data into memory; a *central control unit* to coordinate operations; an *output unit* to record results. This architecture allows one to change the program without changing the machine physically and therefore added a whole new dimension of flexibility. Its ability to modify its own instructions made it 'the first fully-fledged stored-program computer design' (Dyson).

This architecture became known as the 'von Neumann architecture'. If you are reading this as a PDF document on a PC or Mac, then you are using a machine operating with 'von Neumann

---

[316] Ortvay had been trying to get him interested in studying the logical structure of the brain since 1940.

architecture', though von Neumann himself stressed to contemporaries that Turing deserved much of the credit for the original insights. Stan Frankel, head of numerical computation at Los Alamos during the war, said that von Neumann would hand round copies of Turing's 1936 paper and urged its study. 'Many people have acclaimed von Neumann as the 'father of the computer' … but I am sure that he would never have made that mistake himself,' said Frankel. Von Neumann urgently wanted to get these ideas into the public domain and the publication of his paper lifted the veil of wartime secrecy and sparked many new projects.

Over 1945-6, von Neumann successfully lobbied the IAS to allow him to build a state-of-the-art new computer based on the lessons of ENIAC and the EDVAC report.[317] He also encouraged other places to produce clones based on the EDVAC report which meant his eponymous architecture spread (including the MANIAC at Los Alamos[318] and JOHNNIAC at RAND).[319] IBM built a commercial version called the 701. In 1952, Howard Smith used the UNIVAC on TV to predict the outcome of the Presidential election shortly after polls closed. UNIVAC predicted an Eisenhower landslide and this success brought publicity to the idea of computer predictions.

Between the end of the war and his death in 1957, von Neumann wrote a series of papers and lectures concerning the whole range of issues involved in computation and sought to establish a basis for what he called 'a logical theory of automata and information'. This work included, *inter alia*:

(i) *The General and Logical Theory of Automata* (delivered 1948, published 1951).
(ii) Five lectures on *Theory and Organization of Complicated Automata* (delivered in 1951).
(iii) *The Theory of Automata*, a manuscript written 1952–3 but unpublished until after his death.
(iv) *The Computer and the Brain*, drafts for the Sulliman lectures supposed to be delivered to Yale in spring 1956, written 1955 – 6, but unpublished until 1958.

Because of his great workload and cancer, he never got his thoughts into a final published form. The book '*Theory of Self-Reproducing Automata*', published in 1966 (edited by Arthur Burks), comprised edited versions of (ii) and (iii).

Before DNA had been discovered, he was exploring how to construct machines that operate according to natural principles. For example, artificial systems are built on the principle of trying to eliminate error. However, natural systems build themselves with vast numbers of parts some of which inevitably fail yet the effects of local failure are minimised and the system does not fail ('[I]f a living organism is mechanically injured, it has a strong tendency to restore itself. If … we hit a man-made mechanism with a sledge hammer, no such restoring tendency is apparent… It's very likely that on the basis of the philosophy that every error has to be caught, explained, and corrected, a system of the complexity of the living organism would not run for a millisecond… ') How can we

---

[317] Cf. *Preliminary Discussion of the Logical Design of an Electronic Computing Instrument* (1946/7) and the three volume *Planning and Coding Problems for an Electronic Computing Instrument (1947/8)*. The IAS faculty were not enthusiastic about such a practical project arriving in their peaceful idyll, with workmen and hammers, but they knew von Neumann was determined to build it somewhere and they did not want him to leave, particularly since he organised the funding from military sources. Some working on it took advantage of Gödel's lack of a secretary and perched in the allocated office.

[318] George Gamow dubbed it MANIAC: the 'Metropolis And von Neumann Install Awful Computer'.

[319] '[I]n planning anything new … it is customary and very proper to consider what the demand is, what the price is, whether it will be more profitable to do it in a bold way or a more cautious way… Things would very quickly go to pieces if these rules were not observed in 99 cases out of 100. It is very important however that there should be one case in a 100 where it is done differently … to write specifications simply calling for the most advanced machine which is possible in the present state of the art' (Von Neumann). His experience in developing computers led him to believe that after an initial technological breakthrough, competition between alternative pathways was preferable to concentration on one avenue. He advocated this approach for both the hydrogen bomb and the missile programme.

learn from nature, particularly the brain, to make reliable computers from unreliable parts? Can we find ways of using evolutionary principles, which win local battles against entropy, to help us? He invented the concept of cellular automata partly to investigate issues of self-reproduction, error correction, and increasing complexity etc. He gave the first proof that self-reproducing machines of growing complexity are logically possible (extending Turing's idea of a Universal Turing Machine to that of a 'Universal Constructor') and he explored how machines might build themselves using evolutionary concepts such as artificial 'genes' etc. He sought a general logical theory of automata abstracted from natural and artificial systems ('life is a process which can be abstracted away from any particular medium'), the new maths necessary for such a theory[320], and the technology needed to instantiate a new mathematical theory (including copying the combined digital/analog technologies and statistical language of biology).

Since von Neumann, hardware has advanced exponentially with Moore's Law. Tubes were replaced by discrete transistors in the late 1950's; in the mid-1960's microcircuits contained several transistors, then hundreds, thousands and so on. The microprocessor of the early 1970's held a complete computer processor on one silicon chip. The first handheld electronic calculators were produced around 1970.[321] Software has not advanced as much. New methods were developed for writing instructions in symbolic form which the computer would turn into binary. IBM introduced FORTRAN in 1957. In 1964, Kemeny and Kurtz invented BASIC, which brought programming to school kids such as Bill Gates.

Shannon and Information Theory… Claude Shannon had worked on Bush's Differential Analyzer in the 1930s. In 1937, he wrote '*A Symbolic Analysis of Relay and Switching Circuits*' which Gershenfeld calls 'perhaps the most consequential master's thesis ever'. Like Turing, he had done some work on cryptography during the war and after 1945 he got to know Turing who also did some post-war work at Bell.[322] In 1948 he published '*A Mathematical Theory of Communication*' which first appeared in the *Bell System Technical Journal.* This introduced *Information Theory*.[323] Shannon would later write that his work on cryptography and his development of Information Theory 'were so close together you couldn't separate them', particularly because the concept of *redundancy* in a coded message was central to its susceptibility to being decoded.[324]

Gershenfeld explains… The goal had been that a telephone would convert sound waves to electrical waves, send them down a wire, and convert them back to sound waves at the other end. Interference, or 'noise', would inevitably introduce some errors, with errors rising in proportion to

---

[320] 'Everybody who has worked in formal logic will confirm that it is one of the technically most refractory parts of mathematics. The reason for this is that it deals with rigid, all-or-none concepts, and has very little contact with the continuous concept of the real or of the complex number, that is, with mathematical analysis. Yet analysis is the technically most successful and best-elaborated part of mathematics. Thus formal logic is, by the nature of its approach, cut off from the best cultivated portions of mathematics, and forced onto the most difficult part of the mathematical terrain, into combinatorics.' VN.

[321] The microchip was first invented in 1952 by MOD worker Geoffrey Dummer but the MOD refused to fund it and a patent was filed 7 years later by Jack Kilby who worked for Texas Instruments.

[322] He wrote a paper on this in 1945 titled '*A Mathematical Theory of Cryptography*' which was classified; a declassified version was published as '*Communication Theory of Secrecy Systems*' in 1949.

[323] The paper was published with an introduction by Weaver in 1949 as *The Mathematical Theory of Communication.* Also cf. Shannon's *Programming a Computer for Playing Chess* (1950).

[324] When Shannon discussed his ideas with Von Neumann, the latter pointed out that Shannon's work had roots in Boltzmann's observation in his work on statistical mechanics (in the late 19th Century) that *entropy is related to missing information* in the sense that 'it is related to the number of alternatives which remain possible to a physical system after all the macroscopically observable information concerning it has been recorded' (Weaver), and these ideas were extended by Szilard's work in the 1920s on Maxwell's Demon (*On the reduction of entropy in a thermodynamic system by the intervention of intelligent beings*)and by von Neumann in the 1930s in his work on quantum mechanics.

the noise. Engineers would find clever ways of incrementally reducing the amount of errors for a given level of noise. Shannon reassessed the problem fundamentally. He assumed signals sent as binary digits ('bits') and showed that for digital signals, below a certain threshold of noise, the error rate is effectively zero; the imperfect components of a phone system could transmit information perfectly.

> 'For given amounts of noise in the system and energy in the signal, the error threshold corresponds to a communication capacity. If the bits are sent below a certain rate, … they can arrive without errors, and above that they're guaranteed to have errors. The bad news is the existence of the threshold: there is an end point to the engineering advances in the phone system… Shannon's result placed a limit on how much information can be reliably sent. But the threshold is also the good news, because instead of engineering improvements leading to smaller and smaller errors, it showed the qualitatively different possibility of making no errors at all… [S]ignals could be sent anywhere without degradation by converting sounds to and from digital data rather than retaining the analog waves. And the capacity limit has not proved to be much of a limitation at all; it's been pushed up from the thousands to millions to billions and now trillions of bits per second in an optical fiber.' (Gershenfeld)

Shannon also realised that one could add extra information to a signal and use that *redundancy* to identify and correct errors. Von Neumann would apply the same idea to demonstrate fault-tolerant computing.

Linear Programming (LP)

'The development of linear programming is - in my opinion - the most important contribution of the mathematics of the 20th century to the solution of practical problems arising in industry and commerce.' Martin Grötschel.

Dantzig spent the war working on various problems for the military, including scheduling problems such as logistics. At the end of the war, the Pentagon offered him a good deal to work on how to mechanise their planning processes. This work became what is now known as LP. Dantzig described a meeting with von Neumann:

> 'On October 3, 1947, I visited him for the first time at the [IAS]. I remember trying to describe to von Neumann, as I would to an ordinary mortal, the Air Force problem. I began with the formulation of the linear programming model in terms of activities and items, etc. Von Neumann did something which I believe was uncharacteristic of him. 'Get to the point,' he said impatiently. Having at times a somewhat low kindling-point, I said to myself 'O.K., if he wants a quicky, then that's what he will get.' In under one minute I slapped the geometric and algebraic version of the problem on the blackboard. Von Neumann stood up and said 'Oh that!' Then for the next hour and a half, he proceeded to give me a lecture on the mathematical theory of linear programs.

> 'At one point seeing me sitting there with my eyes popping and my mouth open (after I had searched the literature and found nothing), von Neumann said: "I don't want you to think I am pulling all this out of my sleeve at the spur of the moment like a magician. I have just recently completed a book with [Morgenstern] on the theory of games. What I am doing is conjecturing that the two problems are equivalent. The theory that I am outlining for your problem is an analogue to the one we have developed for games."'

After this meeting, the two kept in touch and LP problems were run on the new computers. Dantzig delivered his results to what became a famous meeting in 1948. After he had finished, Hotelling stood up, said 'But we all know the world is nonlinear', and sat down. Fortunately for Dantzig, von Neumann stood up and defended him.

In LP, we 'require the objective to be a linear expression of the variables in the model, and this expression is maximized, or minimized, over all allowable assignments of levels to the activities. An optimal solution to an LP model is an assignment such that the objective is as large, or small, as possible.' Dantzig developed the *simplex algorithm* to analyse LP problems: LP data is fed into the simplex algorithm and out pops an optimal answer. LP enables us to say - 'no tour through this point set can be shorter than X' - therefore if we can find a tour of length X, then we know it is the shortest. LP is the method adopted in all of the most successful exact TSP approaches so far. (Cf. *The Traveling Salesman*, Cook. For a description, and examples of the global use, of LP, cf. *Model Building in Mathematical Programming*, by Paul Williams. Also Bob Vanderbei's web Simplex Pivot Tool.)[325]

The Monte Carlo method

Ulam described how during an illness in 1946 he was playing Solitaire and realised that it was easier to get an idea of the general outcome of the game *via* multiple random trials than a formal computation of all the possibilities. Given his work at the time on problems such as the hydrogen bomb, he saw that the same approach could be applied to a wide range of complex systems then intractable to normal computational methods.

> '… the question was what are the chances that a Canfield solitaire laid out with 52 cards will come out successfully? After spending a lot of time trying to estimate them by pure combinatorial calculations, I wondered whether a more practical method … might not be to lay it out say one hundred times and simply observe and count the number of successful plays. This was already possible to envisage with the beginning of the new era of fast computers, and I immediately thought of problems of neutron diffusion…

> 'The idea was to try out by mental experiment thousands of such possibilities, and at each stage, determine by chance, by a random number distributed according to the given probabilities, what should happen, and total all these possibilities and get an idea of the behavior of the physical process… [The Monte Carlo method] came into concrete form … after I proposed the scheme to Johnny [von Neumann]."

The first computerised Monte Carlo simulation, of a Los Alamos related problem, was carried out on the ENIAC in 1947 (Aspray). Afterwards, the technique spread rapidly and it is used today in many different fields.[326] Cf. This [paper](#) on Ulam, von Neumann, and the 'Monte Carlo method'.

The Internet

'In a universal organization and clarification of knowledge and ideas … in the evocation, that is, of what I have called a World Brain … in that alone is there any clear hope of a really Competent receiver for world affairs… We do not want dictators, we do not want oligarchic parties or class rule, we want a widespread world intelligence conscious of itself… The whole human memory can be, and probably in a short time will

---

[325] 'The simplex algorithm for linear programming has been shown to have exponential time complexity, but it has an impressive record of running quickly in practice' (Garey, p.7ff).

[326] 'To this day, the theoretical statistician is unable to give a proper foundation to the method.' Metropolis, who worked on the Monte Carlo method with Ulam and von Neumann (*The Age of Computing: A Personal Memoir*, 1992).

be, made accessible to every individual… [It] need not be concentrated in any one single place. It need not be vulnerable as a human head or human heart is vulnerable. It can be reproduced exactly and fully in Peru, China, Iceland, Central Africa, or wherever else seems to afford an insurance against danger or interruption. It can have at once the concentration of a craniate animal and the diffused vitality of an amoeba.' HG Wells.

'Probably the closest parallel structure to the Internet is the free market economy.' Baran

It is often stated that the internet was developed to withstand a nuclear attack. This is not quite accurate. Here is a brief history of its genesis.[327]

In the 1950s, Whirlwind evolved into SAGE (Semi-Automatic Ground Environment) to support air defence for the USAF and was the beginning of 'a real-time interactive data-processing system' (Dyson). Von Neumann was involved with SAGE in the 1950's. From 1951 RAND did many experiments on human operators of machines used to defend against air attack. The UCLA students playing the operators improved their performance so well that the USAF asked RAND to train real operators. The USAF contracted with MIT and RAND to develop SAGE after Bell and IBM declined the offer; 'We couldn't imagine where we could absorb 2,000 programmers at IBM when this job would be over someday, which shows how well we were understanding the future at that time,' said Crago.

'The switch from batch processing, when you submit a stack of cards and come back hours or days later to get the results, to real-time computing was sparked by SAGE's demand for instantaneous results … [and] also led the switch to magnetic-core memory.' SAGE ran on about 1 million lines of software code which was 'by far the largest software project of its time' (Dyson).[328] From the late 1950's and early 1960's, various military and civilian projects were working on communication networks, particularly RAND and ARPA (which was formed in 1958 as part of the panicked response to the launch of Sputnik; renamed DARPA in 1972). RAND also experimented with the Leviathan project (from 1959), an early experiment with introducing evolutionary principles into computing rather than having everything hierarchically controlled as with SAGE. The plan was to create artificial agents 'isomorphic with significant aspects of individual human actions' and simulate 'the micro-functioning of real agents under social constraints.' However, the computer power of the time could not produce sufficiently interesting simulations and the project was dropped.[329]

IBM had built 23 computers for the SAGE program which were dispersed in bunkers to interpret radar data. There was an unneeded spare and newly developed ARPA acquired it from the Air Force in 1961.

Licklider's 'Inter-Galactic Computer Network'… Joseph Licklider was a psychologist and very interested in computer science. He had worked on the SAGE program then for BBN, the famous technology firm. On SAGE, he had examined the problem of human-computer interfaces; e.g. how can humans interpret clumsy radar screens. During the 1960's, he wrote a series of memos on what he jokingly called the possibilities for the "Inter-Galactic Computer Network". In 1960, he wrote "*Man-Computer Symbiosis*" and in 1968 "*The Computer as a Communication Device*".[330] He was

---

[327] This section is taken partly from Dyson, RAND and DARPA including Belfiore (2009).

[328] One of the spinoffs from SAGE would be SABRE, the airline ticket reservation system.

[329] Leviathan was itself born of an earlier project developed by Oliver Selfridge called 'Pandemonium' which was aimed at understanding of a Morse code stream sent by human operators using the principles of evolutionary computing to 'evolve' understanding.

[330] Licklider's two most famous papers: http://memex.org/licklider.pdf

hired by ARPA from 1962 to run their computer research program and find uses for their new IBM supercomputer.

In his 1960 paper he wrote:

> 'Man-computer symbiosis is an expected development in cooperative interaction between men and electronic computers… The main aims are 1) to let computers facilitate formulative thinking as they now facilitate the solution of formulated problems, and 2) to enable men and computers to cooperate in making decisions and controlling complex situations without inflexible dependence on predetermined programs… Prerequisites for the achievement of the effective, cooperative association include developments in computer time sharing, in memory components, in memory organization, in programming languages, and in input and output equipment… The hope is that, in not too many years, human brains and computing machines will be coupled together very tightly, and that the resulting partnership will think as no human brain has ever thought and process data in a way not approached by the information-handling machines we know today.

> 'Man-computer symbiosis is probably not the ultimate paradigm for complex technological systems. It seems entirely possible that, in due course, electronic or chemical "machines" will outdo the human brain in most of the functions we now consider exclusively within its province… There will nevertheless be a fairly long interim during which the main intellectual advances will be made by men and computers working together in intimate association. [This interim] may be 10 or 500, but those years should be intellectually the most creative and exciting in the history of mankind.'

There are two main connected goals of the man-machine symbiosis. First, computers now have to be laboriously programmed to do something in completely pre-specified detail, or they crunch, rather than helping man in a collaborative way. Second, we want computers to help human decisions in real time. Progress on both fronts will allow computers to help humans much more effectively with their problems and decisions.

> 'Present-day computers are designed primarily to solve preformulated problems or to process data according to predetermined procedures. The course of the computation may be conditional upon results obtained during the computation, but all the alternatives must be foreseen in advance. (If an unforeseen alternative arises, the whole process comes to a halt and awaits the necessary extension of the program.)…

> 'However, many problems that can be thought through in advance are very difficult to think through in advance. They would be easier to solve, and they could be solved faster, through an intuitively guided trial-and-error procedure in which the computer cooperated, turning up flaws in the reasoning or revealing unexpected turns in the solution. Other problems simply cannot be formulated without computing-machine aid. Poincaré anticipated the frustration of an important group of would-be computer users when he said, "The question is not, 'What is the answer?' The question is, 'What is the question?' " One of the main aims of man-computer symbiosis is to bring the computing machine effectively into the formulative parts of technical problems.

> 'The other main aim is closely related. It is to bring computing machines effectively into processes of thinking that must go on in "real time," time that moves too fast to permit using computers in conventional ways. Imagine trying, for example, to direct a battle with the aid of a computer on such a schedule as this. You formulate your problem today. Tomorrow you spend with a programmer. Next week the computer devotes 5 minutes to assembling your

program and 47 seconds to calculating the answer to your problem. You get a sheet of paper 20 feet long, full of numbers that, instead of providing a final solution, only suggest a tactic that should be explored by simulation. Obviously, the battle would be over before the second step in its planning was begun. To think in interaction with a computer in the same way that you think with a colleague whose competence supplements your own will require much tighter coupling between man and machine than is suggested by the example and than is possible today.'

Licklider analyses the differences between man and machine and their different roles in the intended symbiosis.

'[M]en are noisy, narrow-band devices but their nervous systems have very many parallel and simultaneously active channels. Relative to men, computing machines are very fast and very accurate, but they are constrained to perform only one or a few elementary operations at a time. Men are flexible, capable of "programming themselves contingently" on the basis of newly received information. Computing machines are single-minded, constrained by their " pre-programming." Men naturally speak redundant languages organized around unitary objects and coherent actions and employing 20 to 60 elementary symbols. Computers "naturally" speak nonredundant languages, usually with only two elementary symbols and no inherent appreciation either of unitary objects or of coherent actions.

'Men will set the goals and supply the motivations, of course, at least in the early years. They will formulate hypotheses. They will ask questions. They will think of mechanisms, procedures, and models. They will remember that such-and-such a person did some possibly relevant work on a topic of interest back in 1947, or at any rate shortly after World War II, and they will have an idea in what journals it might have been published. In general, they will make approximate and fallible, but leading, contributions, and they will define criteria and serve as evaluators, judging the contributions of the equipment and guiding the general line of thought. In addition, men will handle the very-low-probability situations when such situations do actually arise. (In current man-machine systems, that is one of the human operator's most important functions. The sum of the probabilities of very-low-probability alternatives is often much too large to neglect.) Men will fill in the gaps, either in the problem solution or in the computer program, when the computer has no mode or routine that is applicable in a particular circumstance.

'The information-processing equipment, for its part, will convert hypotheses into testable models and then test the models against data (which the human operator may designate roughly and identify as relevant when the computer presents them for his approval). The equipment will answer questions. It will simulate the mechanisms and models, carry out the procedures, and display the results to the operator. It will transform data, plot graphs... The equipment will interpolate, extrapolate, and transform. It will convert static equations or logical statements into dynamic models so the human operator can examine their behavior. In general, it will carry out the routinizable, clerical operations that fill the intervals between decisions. In addition, the computer will serve as a statistical-inference, decision-theory, or game-theory machine to make elementary evaluations of suggested courses of action whenever there is enough basis to support a formal statistical analysis. Finally, it will do as much diagnosis, pattern-matching, and relevance-recognizing as it profitably can, but it will accept a clearly secondary status in those areas.'

Licklider describes a future 'time-shared' network.

'It seems reasonable to envision, for a time 10 or 15 years hence, a 'thinking center' that will incorporate the functions of present-day libraries together with anticipated advances in information storage and retrieval and the symbiotic functions suggested earlier in this paper. The picture readily enlarges itself into a network of such centers, connected to one another by wide-band communication lines and to individual users by leased-wire services. In such a system, the speed of the computers would be balanced, and the cost of the gigantic memories and the sophisticated programs would be divided by the number of users.'

He then describes the problems with the *interfaces* between man and machine: 'there is almost no provision for any more effective, immediate man-machine communication than can be achieved with an electric typewriter'. With displays, 'Nowhere, to my knowledge, however, is there anything approaching the flexibility and convenience of the pencil and doodle pad or the chalk and blackboard used by men in technical discussion.'

'... for effective man-computer interaction, it will be necessary for the man and the computer to draw graphs and pictures and to write notes and equations to each other on the same display surface. The man should be able to present a function to the computer, in a rough but rapid fashion, by drawing a graph. The computer should read the man's writing ... and it should immediately post, at the location of each hand-drawn symbol, the corresponding character as interpreted and put into precise type-face. With such an input-output device, the operator would quickly learn to write or print in a manner legible to the machine... He could correct the computer's data, instruct the machine via flow diagrams, and in general interact with it very much as he would with another engineer, except that the "other engineer" would be a precise draftsman, a lightning calculator, a mnemonic wizard, and many other valuable partners all in one.'

He also discusses the idea of controlling computers *via* natural speech particularly in time-critical situations.

'As military system ground environments and control centers grow in capability and complexity, therefore, a real requirement for automatic speech production and recognition in computers seems likely to develop.'

On arriving at ARPA, Licklider immediately began to implement this program across the country. He established contacts between ARPA and other computer centres which he described as 'members of the Intergalactic Computer Network'. One of the connections was with Doug Engelbart at the Stanford Research institute (SRI). Engelbart had written '*Augmented Human Intelligence*' for the Air Force in 1960 and got research funding. Licklider gave Engelbart's group a remote access login to the new ARPA computer in 1963. This helped Engelbart develop ideas about interfaces and gradually more research centres began to share both the ARPA computer and ideas.

In 1968, Engelbart presented his team's ideas at a conference in San Francisco in one of the most famous presentations ever made in the history of computer science. He demonstrated word processing, video chat, real-time collaborative document editing, hypertext links and the mouse. This was revolutionary in a world of teletypes and punch cards and gave people a path to follow that led to the personal computer revolution in the 1970s. One of the announcements in the talk was ARPA's idea of a new 'experimental network'. At the end of his presentation Engelbart later recalled, 'everyone was standing, cheering like crazy.'

Licklider had a big influence on Robert Taylor (who also spent time at RAND).

"There was one other trigger that turned me to the ARPAnet. For each of these three terminals [terminals in his ARPA office connected to different locations], I had three different sets of user commands. So if I was talking online with someone at S.D.C. and I wanted to talk to someone I knew at Berkeley or M.I.T. about this, I had to get up from the S.D.C. terminal, go over and log into the other terminal and get in touch with them.

"I said, Oh man, it's obvious what to do: If you have these three terminals, there ought to be one terminal that goes anywhere you want to go where you have interactive computing. That idea is the ARPAnet."[331]

In 1965-6, Taylor developed the idea. One day he walked down a corridor and explained the idea to the ARPA director; he took 'a million dollars out of the ballistic missile defense budget and put it into my budget right then and there'.

To enable very different and complex computers to speak to each other, ARPA developed the idea of building identical computers that would actually form the network and communicate with their own individual host computer. These 'interface message processors' (IMPs) became known as 'packet switches' because of the system they used, based on work by Kleinrock, Paul Baran, and Donald Davies.

Baran and RAND…[332]

In 1960, RAND published '*Cost of a Hardened, Nationwide Buried Cable Network*' which estimated the cost of 200 hardened facilities at about $1 billion. Paul Baran's job was to figure out how it should be done technically.

He discerned three basic network topologies: (1) *centralised* networks (a star); (2) *decentralised* networks ("a set of stars connected in the form of a larger star"), what we now call a SFN with hubs; and (3) *distributed* networks, (a mesh). The centralised (star) topology was obviously vulnerable to attack. Baran discerned 'with incredible insight' (Barabasi) that the decentralised (SFN hub) model was also vulnerable. He therefore recommended the distributed model.

Baran's first job had been working with the Eckert-Mauchly Computer Company. Baran's first suggestion was to piggyback the AM radio network, which was very decentralised and redundant ('American country music could survive even a worst-case Soviet attack', Dyson), by installing 'a minimal amount of digital logic' so crucial messages could be dispersed. This was rejected on the grounds that the military needed more bandwidth. He therefore decided he would give them 'so much damn communication capacity they won't know what in hell to do with it all.'

Baran described an ideal communication system:

'An *ideal* electrical communications system can be defined as one that permits any person or machine to reliably and instantaneously communicate with any combination of other people or machines, anywhere, anytime, and at zero cost. It should effectively allow the illusion that those in communication with one another are all within the same soundproofed room - and that the door is locked.' Baran (RAND RM-3767-PR, August 1964).

---

[331] 1999 interview here.

[332] Baran and the internet: http://www.rand.org/about/history/baran.html

Baran laid out a plan for 'digital data transmission systems … where computers speak to each other'; since 'there does not seem to be any fundamental technical problem that prohibits the operation of digital communication links at the clock rate of digital computers', we will soon have the hardware we need 'where the intelligence required to switch signals to surviving links is at the link nodes and *not* at one or a few centralized switching centers' (Baran).

The traditional telephone system was a circuit-switched (not packet-switched) system in which 'a central switching authority establishes an unbroken connection for every communication' which is an attractive target for a bomb. A distributed system in which a huge volume of possible connections exists is more robust but it means the intelligence and memory needed also grows. Baran's idea was to distribute this intelligence not only to the switching nodes but to the messages themselves. He called it 'adaptive message block switching' (which Donald Davies (at the UK National Physical Laboratory) shortened to 'packet switching' in 1966). Each message was broken into 1,024 bit strings and given To / From addresses, and information needed to identify how many steps it had taken, which message it was part of, and how to reconstruct it at the end. By looking at the From address and the handover tag, a node could keep track of which other nodes were forwarding messages efficiently. Nodes worked on the 'hot potato' method of throwing a message at a neighbour based on assessment of performance rather than keeping it. (Dyson)

He attributed the idea to Theseus, a mechanical mouse built by Shannon in 1950 that could find its way around a 5x5 maze: 'within the past decade, the research effort devoted to these ends has developed from analyses of how a mechanical mouse might find his way out of a maze, to suggestions of the design of an all-electronic world-wide communications system' (Baran, 1964).

A demonstration on RAND's IBM 7090 showed such a network could be surprisingly resilient to failure and attack. The whole study was published as a RAND report in 1964 (with some volumes classified). The project soon gained wide support in Government. Unfortunately, there were problems with implementing it as originally envisaged; Baran had to advise against going ahead because it was handed to an agency he thought would mangle it. Significant opposition also came from AT&T who saw lucrative monopolies under threat. A senior AT&T executive supposedly told Baran:

> "First, it can't work, and, if it did, damned if we are going to allow the creation of a competition to ourselves."

In the late 1960's, the ideas on packet switching were incorporated into the embryonic ARPANET. The first two nodes of the ARPANET (at UCLA and SRI) communicated on 29 October 1969. SRI received two letters before crashing. It doubled to four nodes that year. In 1971, the first "email" was sent. It reached about 200 nodes, including outside America, by about 1980.[333]

During the 1970s, DARPA worked on 'inter-networking' between different types of communication network, including radio and satellite. This led to TCP - the transmission control protocol - as a common language for the packet-switching IMPs ('routers') to speak to different types of network.

On 27 August 1976, the SRI team put their stuff in a van and drove to the Alpine Beer Garden where they set up their computers on wooden picnic tables. They set up a remote terminal and sent their weekly report to DARPA. This was the beginning of the (van-sized) mobile internet.

---

[333] Interestingly, although the Soviet bloc was officially excluded from the system, Barabasi writes that ingenious local programmers enabled some academics behind the Iron Curtain to access the embryonic internet before 1989.

The ARPANET was *not* a plan to develop a communication network immune to nuclear attack, it was conceived more of a way of linking different groups connected to ARPA to advance general work on computers. It did incorporate Baran's ideas on packet switching but *not* the distributed mesh network model. The original architecture of ARPANET was not intended to be the basis of either email or the web, another example of unintended consequences.

All the histories of computing since 1945 agree that the role of ARPA was enormous and greater than that of any other single institution. In 1975, Minsky (MIT) said that it was 'almost impossible' to name an important computer project that had not originated at ARPA.

One of the most fundamental problems in maths and computer science: P=NP?

A general way of considering P=NP?: is there necessarily a large gap between the time/resources needed to *solve* a problem and the time/resources needed to *verify* a problem has been solved? E.g. Does it *necessarily* take longer to '*solve*' a jigsaw than it does to *verify* that the solution is right?

(1) Gödel showed that no formal system can encode *all and only* the true mathematical statements. (2) Turing showed that 'even if we settle for an incomplete system F, there is *still* no mechanical procedure to sort mathematical statements into the three categories "provable in F", "disprovable in F", and "undecidable in F"' (Aaronson). However, as Gödel first pointed out, this leaves open the question of (3) whether a mathematical statement S has a proof in F with *n* symbols or fewer.

This problem, unlike Hilbert's original problem, is clearly *decidable* and *computable* for in principle we could always just search every possible combination and check if any has a valid *F*-proof of *S*. 'The issue is "merely" that this approach takes an astronomical amount of time: if n=1000 (say), then the universe will have degenerated into black holes and radiation long before a computer can check 2^{1000} proofs'. The question 'can we prove there is no better approach that avoids brute force searches and finds proofs of size n in time polynomial to n?' is the famous 'P=NP?' problem, one of the seven Clay Millennium problems (Aaronson).[334]

The existence of electronic computers almost immediately raised the issue of what we now call *computational complexity*. In 1949, Claude Shannon referred to the problem of how to know if a Boolean circuit is the most efficient possible to realise a given function.[335] In 1953, von Neumann's last paper on game theory concerned how to avoid brute force search in an exponential space for the 'assignment problem'.[336] What is now known as the 'P=NP?' problem was first described in a March 1956 letter from Kurt Gödel to von Neumann as the latter was dying and it lay undiscovered until 1988, meanwhile the problem was reinvented and became famous.

> 'I hope and wish that your condition will soon improve even further and that the latest achievements of medicine may, if possible, effect a complete cure. Since, as I hear, you are feeling stronger now, I would like to take the liberty to write to you about a mathematical problem; your view on it would be of great interest to me.'

---

[334] 'My own view is that P versus NP is manifestly the most important of the seven problems! For if P=NP, then … there is an excellent chance that we could program our computers to solve the other six problems' (Aaronson).

[335] '*The switching of two terminal switching circuits*' (1949).

[336] '*A certain zero-sum two-person game equivalent to the optimal assignment problem*' (1953). Because of pressure on his diary, this 'paper' was actually a partial transcript of a talk.

'Obviously, it is easy to construct a Turing machine that allows us to decide, for each formula *F* of the predicate calculus and for every natural number *n*, if *F* has a proof of length *n*. Let $\Psi(F,n)$ be the number of steps that the machine requires for that and let $\varphi(n) = \max_F \Psi(F,n)$. The question is, how fast does $\varphi(n)$ grow for an optimal machine. One can show that $\varphi(n) \geq Kn$. If there actually were a machine with $\varphi(n) \sim Kn$ (or even only with $\sim Kn^2$), this would have consequences of the greatest magnitude. That is to say, it would clearly indicate that, despite the unsolvability of the *Entscheidungsproblem*, the mental effort of the mathematician in the case of yes-or-no questions could be completely [author's footnote: 'Apart from the postulation of axioms'] replaced by machines. One would indeed have to simply select an *n* so large that, if the machine yields no result, there would then also be no reason to think further about the problem…

'Now it seems to me to be quite within the realm of possibility that $\varphi(n)$ grows that slowly… It would be interesting to know … how significantly *in general* for finitist combinatorial problems the number of steps can be reduced when compared to pure trial and error…'

Efficient algorithms, P and NP…

What do we mean by a 'good' or 'efficient' algorithm that would 'solve' problems like the Travelling Salesman Problem (TSP) in the sense that the algorithm could take any version of the problem and reliably work out the answer? We have to accept that the solution time will grow with the number of cities; the issue is - how fast?

In general, computer scientists regard an 'efficient' algorithm as one that solves a problem of size *n* in a number of steps proportional to n^k for some value *k* where however large *k* is, it is fixed and does not increase as *n* increases. An algorithm of the form n^k is called a '*polynomial time*' algorithm and an algorithm is 'efficient' if its running time can be *upper-bounded* by a *polynomial function* of n. Problems *solvable* by a polynomial time algorithm constitute (roughly defined) the class called **P**, for 'Polynomial time'. In many cases where we have found an algorithm of this type, we have then found ways to bring the value of k to <10. For example, multiplying two n-digit integers can be done in $\sim n^2$ steps using school methods (or faster using advanced methods) and all four basic arithmetical operations are 'polynomial time' processes.

An 'inefficient' algorithm is one whose running time is *lower-bounded* by an *exponential function* of n, such as n! or 2^n. For example, finding prime factors by searching exhaustively takes $\sim 2^n$ steps and the best *proven* method requires $\sim 2^{\sqrt{n}}$ steps.[337] An algorithm of this form is called an '*exponential time*' algorithm. Some problems are proved to require exponential algorithms (this class is called **EXPTIME**). 'Most exponential time algorithms are merely variations on exhaustive search, whereas polynomial time algorithms generally are made possible only through the gain of some deeper insight into the structure of the problem' (Garey).

The class **NP** does *not* stand for non-polynomial time (as some accounts say) but for 'Nondeterministic Polynomial-time' and is 'the class of computational problems for which a solution can be *recognized* in polynomial time, even though a solution might be very hard to

---

[337] A method depending on plausible but unproven conjectures takes $2^{n^{1/3}}$ steps (Aaronson).

find' (Aaronson).[338] The TSP is an example of a problem in the class NP. 'The best result [for TSP] thus far is a solution method, discovered in 1962, that runs in time proportional to $n^2 2^n$. Although not good, this growth rate is much smaller than the total number of tours through n points, which we know is $(n - 1)!/\ 2$' (William Cook). Despite the difficulty, we have solved a 85,900 city version of the TSP in 2006.[339]

This definition of the distinction between efficient/inefficient[340] allows us to ignore most technical questions of the type of technology being used, though there is a vital distinction in some problems between classical and quantum computers: e.g. some problems are known to be solvable in polynomial time on a QC but are not known to be solvable in polynomial time on a classical computer.[341]

This table demonstrates the problem of how quickly exponential-time algorithms become practically useless, assuming a computer capable of executing a billion operations per second.

Running time on a 10^9 operations-per-second computer

	N=10	N=25	N=50	N=100
n^3	0.000001s	0.00002s	0.0001s	0.001s
2^n	0.000001s	0.03s	13 days	40 trillion years

We know that the TSP is in the class NP: we do not know whether it is also in P. The TSP is a particular type of NP-problem known as '*NP-complete*'.[342] An 'NP-complete' problem is a type of NP problem such that, *if* there is an efficient solution for it, *then all* NP problems have an efficient solution, because all NP problems can be mapped onto each other. *That is, finding an efficient algorithm for a single NP-complete problem would prove (and requires) that P=NP, so the 'P=NP?' problem can be reformulated as: 'can any NP-complete problem be solved in polynomial time?'.*

Examples of other NP-complete problems: 'The Process Scheduling Problem' (what's the quickest way to divide up tasks so they can be done as quickly as possible?), the famous problem of colouring a map with three colours without colouring neighbours the same, jigsaws, Sudoku, and Gödel's original problem of finding mathematical theorems. Factoring *is* in NP but is *not* known to

---

[338] A 'Nondeterministic polynomial time' (NP) problem means a problem that can be solved by a non-deterministic Turing Machine that picks the right solution randomly with perfect luck; it is only the number of possibilities that is the problem. It is a theoretical construct and is obviously not realistic. 'There *are* computational problems that are *known* to require more than polynomial time, but the NP problems are not among those. Indeed, the classes NP and 'non-polynomial' have a nonempty intersection exactly if P≠NP' (Aaronson).

[339] The record for solving TSPs rose from 49 (1954), 64 (1971), 80 (1975), breakthroughs in 1987 from 532 to 2,392; 13,509 (1998). A 1,904,711 version of all the world's towns and villages is unsolved, but we know that best solution so far is <0.05% from the optimal result.

[340] Of course, a) an exponential algorithm requiring 1.00000001^n steps would be faster than one requiring n^{1000} steps and b) there are some growth rates that fall between the two categories (like $n^{\log n}$) but in practice the distinction turned out to be very useful (Aaronson).

[341] NB. '… not known to be solvable in polynomial time on a classical computer' does not mean '*known not* to be solvable in polynomial time on a classical computer', a distinction sometimes obscured (see below).

[342] In 1971, Stephen Cook introduced the idea of NP-completeness in '*The Complexity of Theorem Proving Procedures*'. Although he could only prove one abstruse problem was NP-complete, within weeks others proved that other problems were too and the list continued to grow.

be NP-complete and in fact 'there is strong evidence that factoring integers is *not* NP-complete' (Aaronson).[343]

What would be the significance of solving P=NP?

If P=NP, 'then the ability to *check* the solutions to puzzles efficiently would imply the ability to find solutions efficiently' (Aaronson): *'any polynomially verifiable problem would be polynomially decidable' (Sipser)*; if a theorem *has* a proof of reasonable length, then we could *find* that proof in a reasonable time. If P=NP, despite the Gödel and Turing results, Hilbert's dream of mechanizing mathematics would practically prevail. If P=NP, then an algorithm could *solve* problems like the Travelling Salesman as quickly as solutions can be *verified* (imagine being able to *do* a jigsaw as fast as one can *check* whether a 'done' jigsaw has been done right). This would have extraordinary implications for knowledge and practical problems.

> 'Computational complexity theory plays an important role in modern cryptography. The security of the internet, including most financial transactions, depend on complexity-theoretic assumptions such as the difficulty of integer factoring or of breaking DES (the Data Encryption Standard). If P= NP these assumptions are all false. Specifically, an algorithm solving 3-SAT in n^2 steps could be used to factor 200-digit numbers in a few minutes.

> 'Although a practical algorithm for solving an NP-complete problem (showing P=NP) would have devastating consequences for cryptography, it would also have stunning practical consequences of a more positive nature, and not just because of the efficient solutions to the many NP-hard problems important to industry. For example, it would transform mathematics by allowing a computer to find a formal proof of any theorem which has a proof of reasonable length, since formal proofs can easily be recognized in polynomial time. Example theorems may well include all of the [Millennium Prize] problems [such as the Riemann Hypothesis]. Although the formal proofs may not be initially intelligible to humans, the problem of finding intelligible proofs would be reduced to that of finding a recognition algorithm for intelligible proofs. Similar remarks apply to diverse creative human endeavors, such as designing airplane wings, creating physical theories, or even composing music. The question in each case is to what extent an efficient algorithm for recognizing a good result can be found. This is a fundamental problem in artificial intelligence, and one whose solution itself would be aided by the NP-solver by allowing easy testing of recognition theories.' ([Stephen Cook](#))

> 'If [P=NP], then we could quickly find the smallest Boolean circuits that output (say) a table of historical stock market data, or the human genome, or the complete works of Shakespeare. It seems entirely conceivable that, by analyzing these circuits, we could make an easy fortune on Wall Street, or retrace evolution, or even generate Shakespeare's 38[th] play. For broadly speaking, that which we can compress we can understand, and that which we can understand we can predict.' (Aaronson, 2005)

A proof that P=NP would enable optimal solutions to be found to so many intellectual and practical problems that 'revolutionary' would seem an understatement. However, most mathematicians think P≠NP and if so we are doomed to live in a world (as are all possible computers and non-human civilizations) in which we will *never* have efficient methods to solve many important problems (though it is *wrong* to say that 'if P≠NP, then machine intelligence is

---

[343] If P≠NP, 'then there are NP problems that are neither in P nor NP-complete and factoring is one candidate for such a problem' (Aaronson). It is sometimes stated that natural physical systems 'solve' NP problems but this is wrong. This [Aaronson paper](#) explains frequent errors in this topic.

impossible'[344]). While the search for a solution of P=NP? has greatly contributed to the development of the field of computational complexity, the general view is that the problem will require new methods and we may still be very far from a solution.[345]

Computational complexity theory has already affected many other debates, from AI to biology and philosophy (cf. a fascinating essay by a leading theorist, '*Why Philosophers Should Care About Computational Complexity*', Aaronson). For example, computational complexity can illuminate a central argument about evolution. Opponents of evolution sometimes say that Darwinian evolution is as unlikely an explanation for complex adaptation as a tornado assembling a 747 as it passes through a junkyard. However, there is a vital distinction between the randomness of the tornado in a junkyard and the randomness in Darwinian evolution. In principle, a tornado *could* assemble a 747 but 'one would need to wait for an expected number of tornadoes that grew exponentially with the number of pieces of self-assembling junk,' just as in thermodynamics, n gas particles *will* eventually congregate in a corner but only after $\sim c^n$ time, for a constant c (which means it is extremely unlikely to be seen in the real world). In contrast, evolutionary processes can be observed in simulations, and sometimes proved theoretically, 'to find interesting solutions to optimization problems after a number of steps that grows only polynomially with the number of variables' (Aaronson).

A summary of our knowledge of the computational complexity classes (Sipser[346]):

$$P \subseteq NP \subseteq PSPACE = NSPACE \subseteq EXPTIME.$$

1. 'We don't know whether any of these containments [ie. the subset symbol $\subseteq$] is actually an equality.'
2. 'Someone may yet discover a simulation … that merges some of these classes into the same class…'
3. We *can* prove that P $\neq$ EXPTIME 'therefore at least one of the preceding containments is proper, but we are unable to say which. Indeed, most researchers believe that all the containments are proper.' (A 'proper containment' means that, for example, $P \subsetneq NP$, meaning that P would be a subset of NP but is *not equivalent* to it, or 'P is a proper subset of NP'.)

Endnote: Quantum Computation

'*Newton showed that the apple falls to the ground according to the same laws that govern the Moon's orbit of the Earth. And with this he made the old differentiation between earthly and heavenly phenomena*

---

[344] '"P≠NP means computers can never replace human mathematicians" is a forehead-bangingly common misunderstanding. Personally, I see no reason why the brain couldn't be simulated by computer (neuron-by-neuron if necessary), and P≠NP does nothing to challenge that belief. All P≠NP suggests is that, once the robots do overtake us, they won't have a general-purpose way to automate mathematical discovery any more than we do today' (Aaronson).

[345] *Why is P=NP? so hard?* There are three main barriers. '*Relativisation*' rules out diagonalization (crucial to Gödel and Turing) and similar techniques. '*Algebrization* … rules out diagonalization even when combined with the main non-relativizing techniques'. '*Natural proofs*' shows that 'many combinatorial techniques, if they worked, could be turned around to get faster algorithms to distinguish random from pseudo-random functions' (Aaronson). *When will P=NP? be solved?* '… my uncertainty applies not only to the approximate number of years but to the approximate log of that number: decades, centuries, millennia, who knows? Maybe the question should be rephrased: "will humans manage to prove P≠NP before they either kill themselves out or are transcended by superintelligent cyborgs? And if the latter, will the cyborgs be able to prove P≠NP?"' (Aaronson).

[346] Sipser, p.312.

obsolete. Darwin showed that there is no dividing line between man and animal. And Einstein lifted the line dividing space and time. But in our heads, we still draw a dividing line between 'reality' and 'knowledge about reality', in other words between reality and information. And you cannot draw this line. There is no recipe, no process for distinguishing between reality and information. All this thinking and talking about reality is about information, which is why one should not make a distinction in the formulation of laws of nature. Quantum theory, correctly interpreted, is information theory.' Zeilinger, a leader in Quantum Information Theory.

'Ordinary classical information, such as one finds in a book, can be copied at will and is not disturbed by reading it. Quantum information is more like the information in a dream. Trying to describe your dream changes your memory of it, so eventually you forget the dream and remember only what you've said about it. But unlike dreams, quantum information obeys well-known laws.' Bennett

'Since quantum mechanics is correct, why is the planet Mars not all spread out in its orbit?' Fermi to Murray Gell-Mann.

Why Quantum Mechanics? Entanglement and 'teleportation'

Quantum mechanics is famously hard to understand, it can only be properly understood using mathematics beyond what most are taught, and most attempts to explain it to a general audience do not work. However, one of the pioneers of quantum computers, Michael Nielsen, has written a brilliant short essay, '*Why the world needs QM*', that explains, much better than scientists usually do, why we *need* a quantum mechanical theory of the world.

Prelude… In 1935, Einstein proposed an experiment (with Podolsky and Rosen, hence 'the EPR experiment') that was designed to show that QM was incomplete.[347] The EPR paper proposed a thought experiment: imagine two entangled particles separated by a great distance. EPR argued that a) the quantum mechanical view of the world implied that a physical measurement on A would *instantly* affect B (what Einstein called 'spooky action at a distance'); b) this violates any reasonable definition of 'realism'; therefore c) quantum mechanics must be incomplete. Instead, Einstein thought that properties of the entangled particles are somehow determined when they separate, and are 'real' and not affected by measurements (what became known as a 'hidden variables' theory).

In 1964, John Bell came up with a new theory to test the EPR concept. Bell probed the foundations of Einstein's view – *locality* (results of measurements at one location must be independent of anything done simultaneously at another location) and *realism* (the outcome of a measurement on one particle reflects properties that the particle possesses independently from the measurement process). He proved that 'No physical theory of hidden variables can ever reproduce all of the predictions of quantum mechanics' (Bell).

> 'Bell showed that a particular combination of measurements performed on identically prepared pairs of particles would produce a numerical bound (today called a Bell's inequality) that is satisfied by all physical theories that obey these two assumptions [locality and realism]. He also showed, however, that this bound is violated by the predictions of quantum physics for entangled particle pairs.' (Zeilinger)

'Bell's theorem' made it possible to conceive of practical experiments to test the EPR argument. In 1982, Alain Aspect *et al* managed to do a real EPR experiment which vindicated quantum

---

[347] '*Can quantum-mechanical description of reality be considered complete?*', by EPR, *Physical Review*, 1935.

mechanics, though loopholes remained. Since then, more sophisticated experiments have been done (including the 'GHZ experiments') and quantum mechanics has been vindicated every time.

Nielsen explains… The real world, as demonstrated in repeated experiments, does not conform to our 'common sense' assumption that properties of objects have an objective intrinsic existence independent of any 'measurement process' (such a measurement process does not require any conscious individual).

If we think of a tossed coin, we think of whether it has landed heads/tails as an objective property of the object independent of any measurement process; we measure afterwards, by looking at it, what has happened but the fact of being heads/tails is an intrinsic property of the coin.

> 'Photons, like coins, can have binary properties. One of those properties is something called *polarization*. You're probably already familiar with polarization, although you may not realize it. If you take a pair of sunglasess, and hold them up towards the surface of the ocean or a pool on a sunny day, you'll notice that depending on the angle you hold the sunglasses, different amounts of light come through. What this means is that depending on the angle, different numbers of photons are coming through.

> 'Imagine, for example, that you hold the sunglasses horizontally… The photons that make it through the sunglasses have what is called *horizontal polarization*. Not all photons coming toward the sunglasses have this polarization, which is why not all of the photons make it through. In our earlier language, what's going on is that the sunglasses are *measuring* the photons coming toward the sunglasses, to determine whether or not they have horizontal polarization. Those which do, pass through the sunglasses; those which do not, are blocked.'

Think about a tossed coin again.

> 'The usual way we think about the world is that the coin is either heads or tails, and our measurement reveals which. The coin intrinsically "knows" which side is facing up, i.e., its orientation is an intrinsic property of the coin itself. By analogy, you'd expect that a photon knows whether it has horizontal polarization or not. And it should also know whether it has a polarization at 45 degrees to horizontal or not.

> 'It turns out the world isn't that simple. What I'll now prove to you is that there are fundamental physical properties that don't have an independent existence like this. In particular, we'll see that prior to Alice measuring the A or B polarization, the photon itself does not actually know what the value for A or B is going to be. This is utterly unlike our everyday experience – it's as though a coin doesn't decide whether to be heads or tails until we've measured it…

> 'To prove this, what we'll do is first proceed on the assumption that our everyday view of the world is *correct*. That is, we'll assume that photons really do know whether they have horizontal polarization or not, i.e., they have intrinsic values A = 1 or A = -1… We'll find that this assumption leads us to a conclusion that is *contradicted* by real experiments. The only way this could be the case was if our original assumption was in fact *wrong*, i.e., photons don't have intrinsic properties in this way.'

Nielsen then explains a simple experiment measuring the polarisation of photons. On the assumption that the photons have an intrinsic property - polarisation - that 'objectively exists' independently of any 'measurement process', he defines a simple numerical inequality (what's

known as 'the CHSH inequality') that experiments would have to satisfy; the sum of averages of certain properties must be less than or equal to 2. Such experiments have been done (first by Alain Aspect in the 1980s) and found a numerical result of c. 2.8! Further, quantum mechanics predicts precisely this value of 2.8 so not only does the experiment disprove our common sense assumption about the world but it also suggests that conventional quantum mechanics is an accurate theory. All subsequent experiments have found the same phenomenon.

That is, the assumption of objective intrinsic properties of photons existing independently of any measurement process leads to a mathematical requirement of a measurement process that is violated by actual experiments. The only reasonable conclusion is that our assumption of objective intrinsic properties existing independently of any measurement process is wrong.[348]

> 'The analysis done in this essay can be extended to nearly all physical properties. In principle, it holds even for everyday properties like whether a coin is heads or tails, whether a cat is alive or dead, or nearly anything else you care to think of. Although experiments like the Aspect experiment are still far too difficult to do for these much more complex systems, quantum mechanics predicts that in principle it should be possible to do an experiment with these systems where the CHSH inequality fails. Assuming this is the case – and all the evidence points that way – at some fundamental level it is a mistake to think even of everyday properties as having an intrinsic independent existence.'

Contra Einstein, some form of 'spooky action at a distance' does exist. Further, not even abandoning 'locality' in an attempt to save 'realism' works.

> 'Two things are clear from these experiments. First, it is insufficient to give up completely the notion of locality. Second, one has to abandon at least the notion of naïve realism that particles have certain properties ... that are independent of any observation.' (Zeilinger)

However, *this spooky action is not 'faster than light communication'*, as is sometimes reported, and it does not violate Special Relativity, so no classical information can be transmitted using this effect.[349] There is still some argument that not every loophole has been absolutely definitively closed. Probably only space-based experiments will finally close the debate as experiments can be done on entangled particles separated by very large distances such that, due to special relativity, 'both observers can claim that they have performed the measurement on their system prior to the measurement of the other observer. In such an experiment it is not possible anymore to think of any local realistic mechanisms that potentially influence one measurement outcome according to the other one' (Zeilinger).

While we cannot use entanglement directly for instant communication, *we can use this effect for 'quantum teleportation' and 'quantum communication'.*

---

[348] In '*Entanglement Swapping between Photons that have Never Coexisted' (Phys. Rev. Lett., 2013)*, the authors show that the nonlocality of QM means that particles can be entangled *even if they do not exist simultaneously*: 'the nonlocality of quantum mechanics, as manifested by entanglement, does not apply only to particles with spacelike separation, but also to particles with timelike separation. In order to demonstrate these principles, we generated and fully characterized an entangled pair of photons that have never coexisted. Using entanglement swapping between two temporally separated photon pairs, we entangle one photon from the first pair with another photon from the second pair. The first photon was detected even before the other was created. The observed two-photon state demonstrates that entanglement can be shared between timelike separated quantum systems.'

[349] Bell gave an analogy: imagine a man (who actually existed) who always wears a pink sock and a green sock; if one sees one sock, one would know the colour of the other sock without any need for communication between the socks; *it is a logical necessity, not communicated information* (Gell Mann).

Teleportation...

One of the handful of main scientists working on quantum teleportation is Anton Zeilinger. He explains how it works (*Scientific American,* 2004):

> 'Alice and Bob … share an entangled auxiliary pair of photons, Alice taking photon A and Bob photon B. Instead of measuring them, they each store their photon without disturbing the delicate entangled state… Alice has a third photon – call it photon X – that she wants to teleport to Bob. She does not know what photon X's state is, but she wants Bob to have a photon with that same polarization state. She cannot simply measure the photon's polarization and send Bob the result. In general, her measurement result would not be identical to the photon's original state [because of the Uncertainty Principle].

> 'Instead, to teleport photon X, Alice measures it jointly with photon A, without determining their individual polarizations… She might find … that their polarizations are 'perpendicular' to each other (she still does not know the absolute polarization of either one, however). Technically, the joint measurement entangles photon A and photon X and is called a Bell-state measurement. Alice's measurement produces a subtle effect: it changes Bob's photon to correlate with a combination of her measurement result and the state that photon X originally had. In fact, Bob's photon now carries her photon's X's state, either exactly or modified in a simple way.'

Then, to complete the teleportation, Alice sends a conventional signal to Bob telling him the result of her measurement. Depending on the result of Alice's measurement, Bob transforms his photon B in a certain way *which thereby becomes an exact replica of the original photon X.*

> 'Which of the possible results Alice obtains [when measuring A and X together] is completely random and independent of photon X's original state. Bob therefore does not know how to process his photon until he learns the result of Alice's measurement. One can say that Bob's photon instantaneously contains all the information from Alice's original, transported there by quantum mechanics. Yet to know how to read that information, Bob must wait for the classical information, consisting of two bits that can travel no faster than the speed of light… Because a photon's quantum state is its defining characteristic, teleporting its state is completely equivalent to teleporting the particle.'

Teleportation does *not* copy the original. The 'no cloning theory' shows that cloning a quantum state would allow us to violate the Uncertainty Principle therefore it is impossible; teleportation destroys the information in photon X. Heisenberg's Uncertainty Principle has been evaded but not refuted *since teleporting works only if we do not know what we are teleporting.*

> '[P]hoton X's state has been transferred to Bob with neither Alice nor Bob learning anything about what the state is. Alice's measurement result, being random, tells them nothing about the state. This is how the process circumvents Heisenberg's principle, *which stops us from determining the complete quantum state of a particle but does not preclude teleporting the complete state so long as we do not try to see what the state is.* [Emphasis added]

> 'It might seem as if information has traveled instantly from Alice to Bob, beating Einstein's speed limit. Yet this strange feature cannot be used to send information, because Bob has no way of knowing [the state of] his photon… Only when he learns the result of Alice's Bell-state measurement, transmitted to him via classical means, can he exploit the information in the teleported quantum state. If he tries to guess in which cases teleportation was instantly

successful, he will be wrong 75 percent of the time, and he will not know which guesses are correct. If he uses the photons based on such guesses, the results will be the same as they would had he taken a beam of photons with random polarizations. Thus Einstein's relativity prevails; even the spooky instantaneous action at a distance … fails to send usable information faster than the speed of light.'

It is important to realise that 'the objective randomness of the individual quantum event, so much disliked by Einstein, prohibits quantum physics from violating Einstein's own theory of relativity.'

At the quantum level, one can say that swapping the information concerning properties actually amounts to transferring everything relevant – *there is no useful distinction between teleporting the information and teleporting the matter.*

> '… the question is: how do I recognise an original? I maintain: solely through its properties. Matter itself is completely irrelevant. If I swap all my carbon atoms for other carbon atoms, I am still Anton Zeilinger.
>
> 'The only important thing are my properties, and they are based on the order of the atoms – that's what makes me who I am. The atoms are unimportant in themselves. So when we transfer characteristics during teleportation, in this sense we are actually transferring the originals.'

One of the main applications for teleportation is in quantum computers.

Also cf. Nielsen's blog on the foundational problem of QM.

Quantum Computers

One of the leading researchers (Aaronson) has written an article explaining basic issues and his blog tracks the cutting edge of progress (see here for his MIT lecture notes). The author of the standard advanced textbook, Michael Nielsen (whose work on 'Network Science' is discussed extensively above), has also written an article for the layman explaining the basics ('Quantum computing for everyone'), and has posted a series of YouTube videos 'Quantum Computing for the Determined'.

The origin of discussions about 'quantum computers' is often traced to a talk by Feynman in 1981 written up as '*Simulating Physics With Computers*' (1982). To specify the state of 300 particles, one needs 2^{300} complex numbers called amplitudes, one for each possible outcome of measuring the particles. The fact that nature seems to use such vast computational resources to keep track of small numbers of physical objects drew Feynman to the field of quantum computation.[350]

> 'During the 1960s and 1970s, Richard Feynman was involved in attempts to use classical digital computers to evaluate the consequences of quantum field theory. He observed that quantum mechanics was hard to program on a classical digital computer. The reason for this difficulty was straightforward: Quantum mechanics possesses a variety of strange and counterintuitive features, and features that are hard for human beings to comprehend are also hard for classical computers to represent at the level of individual classical bits. Consider

---

[350] In the classical world, if we don't know the value of an n-bit string, then we also describe our ignorance using exponentially many numbers, but the string of numbers really has a definite value and the vector of 2^n probabilities is just a mental representation of our own ignorance. 'With a quantum state, we do not have the same luxury because of the phenomenon of *interference* between positive and negative amplitudes' (Aaronson).

that a relatively small quantum system consisting of a collection of 300 electron spins "lives" in $2^{300} \approx 10^{90}$ dimensional space. As a result, merely writing down the quantum state of the spins in a classical form as a sequence of bits would require a computer the size of the universe, and to compute the evolution of that state in time would require a computer much larger than that.

'In 1982, Feynman noted that if one has access to a quantum-mechanical device for representing the state of the spins and for transforming that state, rather than a classical device, then the computation of the time evolution of such a system can be much more economical. Consider a collection of 300 two-level quantum systems, or qubits, one for each electron spin. Suppose that one sets up or programs the interactions between those qubits to mimic the dynamics of the collection of spins. The resulting device, which Feynman called a universal quantum simulator, will then behave as a quantum analog computer, whose dynamics form an analog of the spin dynamics.

'Since Feynman's proposal, researchers in quantum information have created detailed protocols for programming quantum analog computers, including reproducing the behavior of fermions and gauge fields. Large-scale quantum simulators have actually been constructed out of crystals of calcium fluoride. Each crystal contains a billion billion spins, which can be programmed using techniques of nuclear magnetic resonance to simulate the behavior of a wide variety of solid-state systems. These solid-state quantum simulators have already revealed a variety of previously unknown quantum phenomena, including spin transport rates that are startlingly higher than the rates predicted by semiclassical theory.' ([Lloyd, 2008](#)).

The standard media account of how QCs work is wrong:

'In that account, quantum computers work by exploiting what is called "quantum parallelism". The idea is that a quantum computer can simultaneously explore many possible solutions to a problem. Implicitly, such accounts promise that it's then possible to pick out the correct solution to the problem, and that it's this which makes quantum computers tick. Quantum parallelism is an appealing story, but it's misleading. The problem comes in … picking out the correct solution. Most of the time this turns out to be impossible… [I]n some cases you can mathematically prove it's impossible. In fact, the problem of figuring out how to extract the solution … is really the key problem.' (Nielsen).

The power of a quantum computer is a flipside of the way in which a quantum system requires more information to describe it than a classical system. An n-bit classical computer requires n numbers to describe it but a QC requires 2^n numbers. A QC involves a 'quantum algorithm' manipulating the qubits (quantum bits), perhaps by shining a laser on them. However, when an algorithm changes just one qubit, the complexity explodes.

'… even though we've manipulated just one qubit, the way we describe the quantum computer changes in an extremely complicated way. It's bizarre: by manipulating just a single physical object, we reshuffle and recombine the entire list of 2^n numbers… It's this reshuffling and recombination of all 2^n numbers that is the heart of the matter. Imagine we were trying to simulate what's going on inside the quantum computer using a conventional computer. The obvious way to do this is to track the way the numbers s1, s2,… change as the quantum computation progresses. The problem with doing this is that even a single quantum gate can involve changes to all 2^n different numbers. Even when n is quite modest, 2^n can be enormous. For example, when n = 300, 2^n is larger than the number of atoms in the Universe. It's just not feasible to track this many numbers on a conventional

computer.' (Nielsen)

This is why it is thought to be infeasible for a conventional computer to simulate efficiently a quantum computer. However, 'What's really clever, and not so obvious, is that we can turn this around, and use the quantum manipulations of all these exponentially many numbers to solve interesting computational problems.'

How does it do this?

> 'When you measure a quantum computer's output, you see just a single, random answer — not a listing of all possible answers. Of course, if you had merely wanted a random answer, you could have picked one yourself, with much less trouble. Thus, the sole reason to prefer a quantum computer is that the subatomic world obeys different laws of probability than the ones we are used to. In everyday life, it would be silly to speak of a "minus 30 percent chance of rain tomorrow," much less a "square root of minus 1 percent chance." However, quantum mechanics is based on numbers called amplitudes, which are closely related to probabilities but can also be negative (in fact, they are complex numbers). Crucially, if an event (say, a photon hitting a screen) can happen one way with positive amplitude, and a different way with negative amplitude, then the two amplitudes can "interfere destructively" and cancel each other out, so that the event never happens at all. The goal in quantum computing is to choreograph a computation so that the amplitudes leading to wrong answers cancel each other out, while the amplitudes leading to right answers reinforce… So far we have only been able to do this in a few areas. The main technical problem is that maintaining the quantum state inside the computer is hard - interference from the outside world 'collapses' the fragile quantum state. The technological problems with scaling a QC to a level that it is useful while limiting such interference are what holds QCs back for now.' (Aaronson, NYT)

For a classical computer, factoring *appears* to be (and may be) exponentially harder than multiplication. In 1994, Shor discovered a quantum algorithm that makes factoring about as easy as multiplication (potentially breaking public key cryptography which the internet depends on). In 1996, Grover showed that, while a classical computer has to search exhaustively through all the records of an unsorted database and on average will have to search half the records, a QC could find the result in only the *square root* of the number of records (in general, 'the number of quantum searches required to locate what you're looking for is the square root of the number of places in which it could be', Lloyd), so it could speed up many optimisation problems (though only quadratically, not exponentially).[351] It is also possible to map optimisation problems onto quantum systems and use the quantum computation of the physical system to 'solve' the optimisation problem:

> 'Many classically hard problems take the form of optimization problems - for example, the traveling salesman problem... Such optimization problems can be mapped onto a physical system, in which the function to be optimized is mapped onto the energy function of the system. The ground state of the physical system then represents a solution to the optimization problem' (Lloyd, op cit).

> 'Many interesting but practically intractable problems can be reduced to that of finding the ground state [the lowest energy configuration] of a system of interacting spins; however, finding such a ground state remains computationally difficult. [The most studied example is the 'Ising spin model', where spins may take one of two values (up/down) along a preferred axis.] It is believed that the ground state of some naturally occurring spin systems can be

---

[351] Grover's algorithm has also been shown to be optimal.

effectively attained through a process called quantum annealing. If it could be harnessed, quantum annealing might improve on known methods for solving certain types of problem... ['To implement a processor that uses quantum annealing to help solve difficult problems, we would need a programmable quantum spin system in which we could control individual spins and their couplings, perform quantum annealing and then determine the state of each spin.'] Here we use quantum annealing to find the ground state of an artificial Ising spin system... [S]uch a system may provide a practical physical means to implement a quantum algorithm, possibly allowing more-effective approaches to solving certain classes of hard combinatorial optimization problems… Although our manufactured spin system is not yet a universal quantum computer, by adding a new type of coupler between the qubits, universal quantum computation would become possible.' ([Johnson *et al*, Nature May 2011](#)).

However, it is also often wrongly claimed that we know that a QC could solve NP-complete problems. For example, *The Economist* (Orion's Belt, 2/07) wrote that QCs in principle will mean that 'the optimal solution to a given NP-complete problem can be found in one shot.' Wrong: it is *possible* that we will discover that a computer could solve NP-complete problems efficiently but QCs are neither known nor generally believed to be able to do so. Further, despite Shor's algorithm which shows a particular quantum algorithm that is faster than any existing classical algorithm, we have not *proved* that QCs can solve certain problems faster than classical computers.[352]

These and other discoveries (such as [the innovation of 'quantum game theory'](#) using quantum algorithms to implement strategies), and the potential for quantum information theory to be used to guarantee secure communications, has particularly increased interest in QCs among intelligence services. Perhaps the most important application of QCs will be that originally described by Feynman - the ability to simulate accurately quantum physics and chemistry.

> 'Since Feynman's proposal, … large-scale quantum simulators have actually been constructed… These solid-state quantum simulators have already revealed a variety of previously unknown quantum phenomena... Quantum information theory has shown that, far from being exotic, entanglement [which Einstein famously called 'spooky action at a distance'] is ubiquitous. Entanglement underlies the stability of the covalent bond; entanglement is a key feature of ground states of solid-state systems; even the vacuum of space is entangled! … As its relationship with the theory of quantum matter becomes more elaborate and more intimate, quantum information theory has the potential to unravel some of the deep mysteries of physics... Meanwhile, quantum information theory has suggested a wide variety of fundamental experiments in quantum matter... Unexpected connections have arisen between transport theory in quantum matter and quantum information theory: Hopping electrons might be able to discern winning strategies for chess or Go more efficiently than classical chess or Go masters, or the efficiency of photosynthesis might arise because excitons moving through photocenters effectively implement a quantum computer algorithm.' (Lloyd)

---
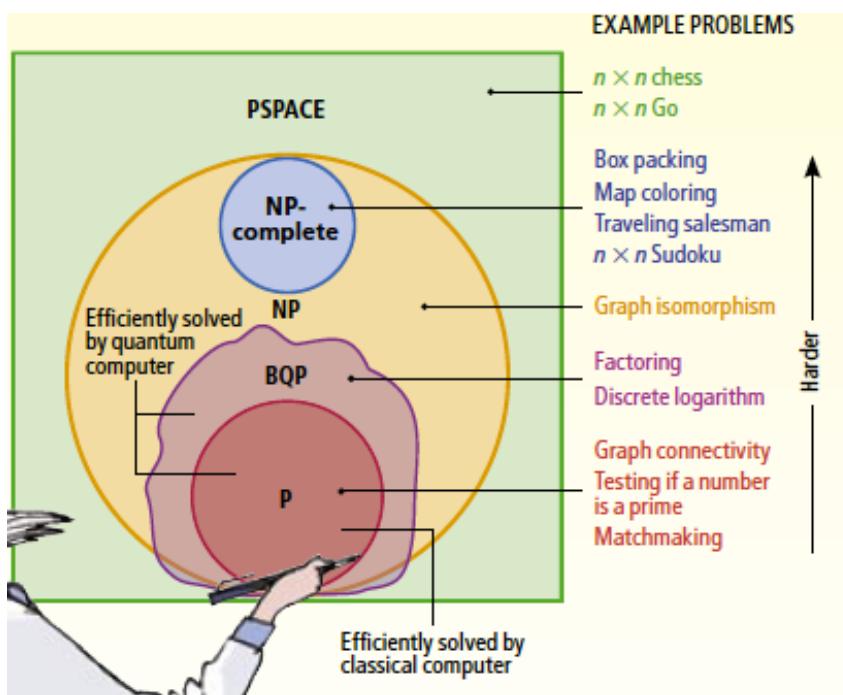
[352] An example of the profound links between maths, computation, and physics: 'It has been suggested that under some circumstances the superposition principle of quantum mechanics might be violated - that is, that the time evolution of quantum systems might be (slightly) nonlinear. Such non-linearity is purely hypothetical: all known experiments confirm the linearity of quantum mechanics to a high degree of accuracy... Nevertheless, the implications of nonlinear quantum mechanics on the theory of computation are profound. In particular, we show that it is generally possible to exploit nonlinear time evolution so that the classes of problems NP and #P ... may be solved in polynomial time. An experimental question - that is, the exact linearity of quantum mechanics - could thereby determine the answer to what may have previously appeared to be a purely mathematical one.' ('*Nonlinear quantum mechanics implies polynomial-time solution for NP-complete and #P problems*', Seth Lloyd, 1998).

One of the basic problems is that quantum systems can 'decohere' on interaction with their environment and lose their 'quantum-ness'. However, ways have been found round this with 'quantum fault-tolerance':

> '… in the 1990s computer scientists and physicists developed the beautiful theory of quantum fault-tolerance, which shows that, as long as you can get the decoherence rate below a certain finite threshold (which current estimates put at somewhere around a 10^{-3} probability of failure per qubit per gate time), you can fix the errors caused by decoherence faster than you introduce new ones, and thereby perform an arbitrarily long quantum computation… So, today there are large research efforts in two directions: on the experimental side, to push down the decoherence rate, and on the theoretical side, to find error-correction schemes that can cope with higher decoherence rates.' (Aaronson)

Special issues of _Nature Physics_ (2012) and of _Science_ (2013) summarise the state of the art in quantum information and communication Bennett's original 1993 paper on quantum teleportation.

Diagram of P, NP and QCs (Aaronson)…

The Church-Turing-Deutsch principle: computation and quantum physics…

After Turing's breakthrough it was soon shown that his system (Universal Turing Machines) and Church's system (λ-calculus) were equivalent. Over the next few decades, it was further proved that many other systems were also computationally equivalent to Universal Turing Machines in the sense that each system can implement all computations that we can conceive of on another system. Any such system that is proved to be computationally equivalent to a Universal Turing Machine is called 'Turing complete': eg. Von Neumann's cellular automata, Conway's _Game of Life_, and even a simple system of bouncing billiard balls are all 'Turing complete'. The standard electronic computer is a particular form of the Universal Turing Machine.

These developments led to an idea - 'the Church-Turing thesis': _anything that can be computed can be computed on a TM._ In 1985, after the famous talk by Feynman in 1981 which introduced the idea

of a quantum computer, [David Deutsch extended the idea](#) to what is now known as 'the Church-Turing-Deutsch thesis': *anything that can be computed can be computed on a quantum TM*. A development of this idea is the 'Church-Turing-Deutsch principle' (CTDP): *any physical process can be simulated on a quantum TM*.

> 'Every finitely realizable physical system can be perfectly simulated by a universal model computing machine operating by finite means…
>
> 'The statement of [this] principle is stronger than what is strictly necessitated by [the original Church-Turing thesis]. Indeed it is so strong that it is not satisfied by Turing's machine in classical physics. Owing to the continuity of classical dynamics, the possible states of a classical system necessarily form a continuum. Yet there are only countably many ways of preparing a finite input for T. Consequently T cannot perfectly simulate any classical dynamical system… I shall show that it is consistent with our present knowledge of the interactions present in Nature that every real (dissipative) finite physical system can be perfectly simulated by the universal quantum computer Q. Thus quantum theory is compatible with the strong form of the Church-Turing principle.'

The CTDP is not proven but there are no known exceptions and most scientists believe it. It is remarkable that not only can a single type of physical object (a physical version of a quantum UTM) perform *any* possible computation that any other physical object could perform, but that it might also be able to simulate *any* physical process. (Cf. Nielsen's short essay on these issues [here](#).)

Endnote: Prizes

Evolution works with lots of small variations and some big leaps on the fitness landscape. However, companies evolve to find 'local maxima'; they do not necessarily make a big leap, hence it often takes a new company to make a breakthrough and this causes established firms to die. The trial and error in an evolutionary variation/selection/failure model, such as the Silicon Valley system of venture capital, has been very successful. Sometimes other things are needed to bring breakthroughs. The Manhattan Project is generally not a good example, being both extremely expensive and secret.
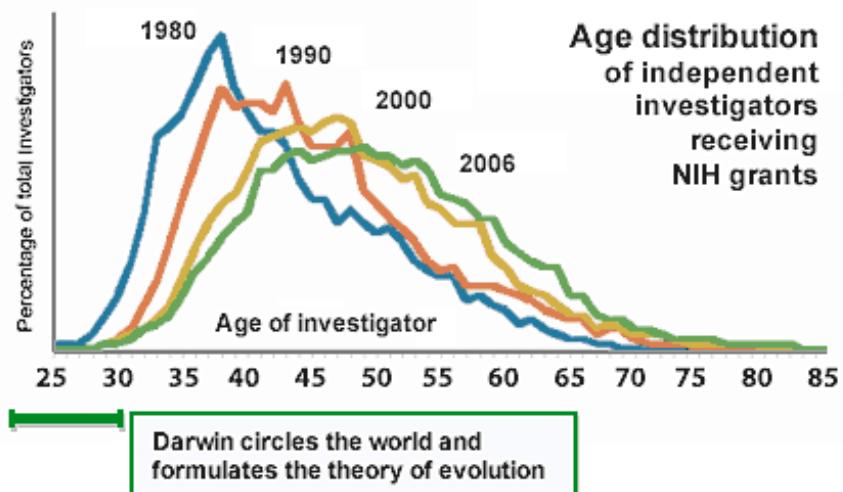
The prize or 'Grand Challenge' approach has worked and is an inspiration for initiatives such as The Gates Foundation's large prizes for drug breakthroughs…E.g. In 1675 the Royal Observatory was founded with the goal of solving the longitude problem. After thirty years little progress had been made. After a particularly bad disaster, the Act of Longitude was passed in 1714 with a prize of £20,000. In 1737, John Harrison, a village carpenter, presented his clock able to keep time at sea.

In 1931 the British Air Ministry issued a specification for a new fighter with extreme performance. Conventional wisdom had been that bombers will always get through. Fighters were not a priority and were built on a conventional biplane model that had evolved in World War I. Sir John Salmond, Chief of the Air Staff, was advised that the RAF faced a big problem: companies were pursuing standard designs which they had 'accumulated a great deal of knowledge' about and involved 'a minimum of unknowns', and although the deficiencies of the standard biplane were becoming obvious and progress depended on 'novel types', 'there can be no chance of developing these unless experience is gained in adapting novel ideas to Service specifications'. Salmond wrote to Dowding in 1931: 'firms are reluctant to risk their money on highly speculative ventures of novel design. If we are to get serious attempts at novel types to meet this specification, we shall have to provide the incentive.' Supermarine did not have the best official entry but Air Commodore Cave went for them and put up £10,000 for the prototype. What emerged was the Spitfire.

The Orteig Prize of $25,000 for the first non-stop Atlantic flight, put up in 1919, inspired nine teams to spend $400,000 to try to win and within three years of Lindbergh's success there was a 30-fold increase in passenger air traffic. This success inspired Diamandis to offer what became known as the Ansari X Prize for the first repeatable private suborbital flight. That $10 million prize inspired $100 million in investment by competitors and success came in 2004. Diamandis extended his X Prize idea to other fields (for example, a $1.2 million prize produced a four-fold improvement in oil spill clean-up technology after the recent BP disaster). With Google, he has set up the Google Lunar X Prize, a $30 million prize for the first privately funded team to send a robot to the moon and have it travel 500 metres and send back data by 2015. (*Wired* interview with Diamandis)[353]

Knowing what to fund is hard. Mario Capechi, a protege of Watson at Harvard, left Harvard for Utah to have space for new projects. In 1980 he sought funding for two 'standard' projects and a very speculative project. NIH agreed to the first two but not the third. He used the funds for the third, risking disaster with NIH. He made a breakthrough and in 2007 won the Nobel. The NIH process (small steps) is designed to be expert-led, rational, and reduce risks: a clear project; experts think it's sound; there's good reason to think it could work. It makes it hard to fund risky projects. The HHMI process (big leaps) funds riskier projects and allows much more freedom. A study comparing the two has found that the HHMI money produced more valuable science and more failures. The HHMI accounts for a twentieth of 1% of global R&D. (Harford, *Adapt*)

Graph: The age distribution of independent investigators receiving NIH grants (source, Drexler).

Answer from the hero in Leo Szilard's 1948 story "The Mark Gable Foundation" when asked by a wealthy entrepreneur who believes that science has progressed too quickly, what he should do to retard this progress: 'You could set up a foundation with an annual endowment of thirty million dollars. Research workers in need of funds could apply for grants, if they could make a convincing case. Have ten committees, each composed of twelve scientists, appointed to pass on these applications. Take the most active scientists out of the laboratory and make them members of these committees… First of all, the best scientists would be removed from their laboratories and kept busy on committees passing on applications for funds. Secondly the scientific workers in need of funds would concentrate on problems which were considered promising and were pretty certain to lead to publishable results… By going after the obvious,

---

[353] NASA has used prizes: e.g. the Strong Tether Challenge, for a material 50% stronger than anything commercially available; the Power Beaming Challenge; the Nanosatellite launch challenge; the sample return robot challenge.

pretty soon science would dry out. Science would become something like a parlor game… There would be fashions. Those who followed the fashions would get grants. Those who wouldn't would not.'

Endnote: Problems with science

Various scientists are noticing a decline in their ability to replicate previously striking results. Eg. Schooler published research in 1990 on how describing an image reduces one's ability to remember it ('verbal overshadowing') but has spent a lot of time over the past 20 years worrying about the decreasing results from follow-up studies. (Similarly, cf. Moller's work on symmetry and Darwinian success.)

An obvious answer is regression to the mean but many of the data sets are big enough that it should almost never happen. Are scientific publishers biased against null findings? Yes, especially where there are incentives not to publicise null results.

When a large number of studies have been done on the same subject, those studies with a large sample size should cluster around a common value while those with a small sample randomly scatter, and this distribution gives the name 'funnel graph'. *Various funnel graph analyses have shown that studies do not distribute randomly but are skewed towards positive results.* E.g. Studies of acupuncture have shown very different results in Asia and Europe / North America, suggesting that scientists unconsciously bias measurements to support prior beliefs.

In 2005, Ioannidis published a controversial paper titled '*Why most published research findings are false*.' He found that 80% of non-randomized studies (by far the most common type) are wrong. Ioannidis examined the 49 most cited clinical-research studies (mostly 'gold standard' randomised trials) in three major journals and found that 45 reported positive results suggesting that the intervention being tested was effective. Of the 34 that had been retested, between a third and half the claims were wrong; 11 had not even been retested. Ioannidis says scientists are chasing the magic 95% boundary and therefore pull data around to make it show something. 'It feels good to validate a hypothesis. It feels even better when you've got a financial interest in the idea or your career depends upon it.' Ioannidis cites poor research (misuse of statistical software), financial conflicts of interest (drug companies, incentives to be published/funded, 'There is an intellectual conflict of interest that pressures researchers to find whatever it is that is most likely to get them funded'), the tendency to focus on exciting theories, patients misreport diet, studies too short to track to death etc. 'Even when the evidence shows that a particular research idea is wrong, if you have thousands of scientists who have invested their careers in it, they'll continue to publish papers on it. It's like an epidemic, in the sense that they're infected with these wrong ideas, and they're spreading it to other researchers through journals… Some fear that there may be less funding… But if we can't really provide those miracles, how long will we be able to fool the public anyway?'

A follow-up study found that ~⅔ of attempts by Bayer to confirm published claims failed. A major meta-meta-study ('*Quality of Systematic Reviews of Observational Nontherapeutic Studies*') showed profound problems with many biological / medical research studies (2010). Nature also reported recently on failure to replicate 47/53 'landmark' studies.

Endnote: intelligence, IQ, genetics, and extreme abilities

'The stars above us govern our conditions; else one self mate and mate could not beget such different issues.' Lear.

This note comes mainly from Robert Plomin's textbook '*Behavioral Genetics*' and blogs and talks by Steven Hsu, Professor of Physics at the University of Michigan. Hsu is an unusual professor. He has been a tech start-up entrepreneur and is extremely interested in IQ and genetics. He is also now helping BGI's project to explore the genetic basis of IQ with a 'genome wide association study' (GWAS). His great blog has lots of relevant material and this fascinating 2013 talk explains a lot. The summary below is an attempt at summarising a lot of mainstream scientific material.

A summary… IQ ≈ 'general cognitive ability' ('g'); 'g' *can* be statistically identified; it *is* heritable; it is *not* measuring the income of parents; high 'g' *is* required for breakthroughs in maths and physics; 'g' predicts key social outcomes such as educational and occupational levels far better than any other trait' (Plomin, 2013); while 'g' is important, obviously other things also contribute to achievements and failures.

NB. Many of these statements are statistical. The fact that, for example, it is very *unlikely* for someone below the 90th percentile to be able to do a maths PhD does not mean it is *impossible* for an individual. Outliers are important.

What are 'g' and 'IQ'?

Different types of test measure different cognitive abilities (e.g. verbal ability, maths ability, spatial ability, memory, speed of processing). 'These broad factors intercorrelate modestly. In general, people who do well on tests of verbal ability tend to do well on tests of spatial ability.' *'g' is that which is in common among these broad factors.* IQ tests generally combine different specific tests into a score that estimates 'g': ie. IQ ≈ 'g'.

> 'In research contexts, g is usually derived by using a technique called *factor analysis* that weighs tests differently according to how much they contribute to g. This weight can be thought of as the average of a test's correlations with every other test. This is not merely a statistical abstraction - one can simply look at a matrix of correlations among such measures and see that all the tests intercorrelate positively and that some measures (such as spatial and verbal ability) intercorrelate more highly than do other measures (such as nonverbal memory tests). A test's contribution to g is related to the complexity of the cognitive operations it assesses. More complex cognitive processes such as abstract reasoning are better indices of g than less complex cognitive processes such as simple sensory discriminations.[354] Although g explains about 40 percent of the variance among such tests, most of the variance of specific tests is independent of g. Clearly there is more to cognition than g… In addition, just as there is more to cognition than g, there is clearly much more to achievement than cognition. Personality, motivation, and creativity all play a part in how well someone does in life. However, it makes little sense to stretch a word like *intelligence* to include all aspects of achievement, such as emotional sensitivity and musical ability that do not correlate with tests of cognitive ability…

> 'There is a wide gap between what laypeople (including scientists in other fields) believe and what experts believe. Most notably, laypeople often read in the popular press that the assessment of intelligence is circular - intelligence is what intelligence tests assess. On the contrary, g is one of the most reliable and valid measures in the behavioral domain. Its long-term stability after childhood is greater than the stability of any other behavioral trait… It is less clear what g is and whether g is due to a single general process, such as executive

---

[354] All existing IQ tests are inevitably 'noisy', not perfect, and are probably only accurate to ~1SD though one can have subjects take more tests to improve accuracy. A single test of a 5 year-old is not a good predictor of adult IQ but the average of multiple tests at 5-7 is accurate to ~0.85 (Hsu).

function or speed of information processing, or whether it represents a concatenation of more specific cognitive processes…' (Plomin 2013, p. 187.)

Is 'g' / IQ heritable?

'[T]here is considerable consensus among scientists … that *g* is substantially heritable.' Heritability is usually measured on a scale between 0 and 1 where 0 means none of the variance in the population is due to genes and 1 means all of it is. NB. This is a population statistic - it does *not* mean that for every individual x% of one's IQ score is accounted for by genes.[355]

A complication in stating a figure for the heritability of *g* is that the picture changes as children age. It seems reasonable to think that genetic differences become *less* important as children age but this is wrong. We have discovered that 'genetic factors become increasingly important for *g* throughout an individual's life span.' (Perhaps different genes affect *g* in adulthood. It is more likely that 'relatively small genetic effects early in life snowball during development, creating larger and larger phenotypic effects', for example as the influence of parents and teachers fade and adults increasingly control their own lives.[356]) Further, studies of adoptive siblings show essentially *zero correlation* when retested as adults; '*shared environment effects' decrease with age.* For teens, ~50% of the variance is due to genes (up from ~20% for small children), ~25% is due to shared environment, ~20% is due to nonshared environment, and ~5% is due to measurement error. *In summary, for adults, ~60% of the variance (at least, increasing towards 70-80% in middle age) is due to genes, ~35% is due to nonshared environment, ~5% is due to measurement error, and ~0% is due to shared environment.* (Plomin 2013.)[357]

> 'One reason to think that heritability estimates are roughly accurate is that the basic quantitative genetic designs used to estimate heritability - family, adoption and twin designs - generally converge on similar estimates of heritability. Each of the designs has potential problems, but they have different problems, which makes this convergence reassuring... Moreover, GWA data generally confirm the basic premise of the twin method...' (Plomin, '*Child Development and Molecular Genetics*', 2013).

Plomin says that the GWAS are the 'beginning of the end' of the long argument about 'nature v nurture' because 'it is much more difficult to dispute results based on DNA data than it is to quibble about twin and adoptee studies'. In 2011, a GWAS confirmed the rough numbers from the twin/adoption studies for IQ ('*Genome-wide association studies establish that human intelligence is highly heritable and polygenic*', *Nature*, October 2011).

Although it is often claimed that current knowledge about the complexity of gene-gene interactions etc shows 'we can't really understand how genes affect IQ', this is false; we can separate out the *linear* effects ('additive heritability') from the *nonlinear* effects and the former accounts for most of the heritability (i.e most of the heritability is linear, not caused by complex

---

[355] Twin correlations for g are ~0.85 for identical twins and ~0.60 for fraternal twins.

[356] 'The authors find a substantial genetic correlation (0.62) between intelligence in childhood and in old age, which means that many of the same genetic elements are associated with this trait throughout life. The analysis also estimates the genetic influence on cognitive change across life: nearly a quarter of the variation in the change in cognitive scores that occurs throughout life could be explained by different genes being associated with this trait in childhood and later life.' Plomin, *Nature*, 2012.

[357] There is also the issue of 'assortative mating' - the tendency of people to have children with similar people. 'Assortative mating for *g* is substantial, with average spousal correlations of about 0.40. In part, spouses select each other for *g* on the basis of education. Spouses correlate about 0.60 for education, which correlates about 0.60 with *g*.' This *increases genetic variance in the population.* How quickly is it changing society?

gene-gene interactions). Recently, some discoveries concerning *epigenetics* have been distorted in various media and popular science books to confuse the picture on 'g' and heritability.

Tests, IQ, effort, and achievement

In America, SATs ≈ IQ. Despite shortcomings that make them 'noisy', SATs and similar tests can predict remarkably well future performance. (SATs are more like IQ tests than GCSEs and some institutions introduced them because some US states banned IQ tests.) Although there are outliers, generally it is hard to improve scores on SATs by more than a few points even with intense effort and private tuition. Enrichment is 'unlikely to drastically alter cognitive ability. An 1150 SAT kid is not going to become a 1460 or 1600 kid as a result of their [college education].' (Hsu) There has been a lot of evidence that specific training programmes can improve specific skills but little, and very controversial evidence, that such improvements *transfer to untrained tasks*. A paper by Jaeggi & Buschkuehl (Proceedings of the National Academy of Sciences, 5/2008) claimed that training working memory improved performance in a *fluid intelligence test*. This got a lot of attention (e.g. from the Office of Naval Research) but many have criticised it and it has not yet been reliably replicated.

According to research by Hsu, hard work *can* yield high GPAs in humanities courses. However, maths and physics are different. Here there appears to be a 'hard threshold'; students below the 90[th] percentile in ability in maths SATs almost never do well enough to get on a PhD programme *no matter how hard they work*.[358]

The arguments that 'IQ above 120 doesn't matter' and 'there's no link between cognitive ability and creativity' are flawed. The best studies we have (Terman, Roe, 'Study of Mathematically Precocious Youth') show that high IQ is important for scientific success, including 'creativity', and we can identify such people.

While the typical science doctoral student has an IQ about 2SDs above the norm (i.e. ~130), eminent scientists have IQs significantly beyond this; +3-4SD or higher (Roe). If IQ above 120 were 'irrelevant', as Gladwell and others claim, then we would expect to see eminent scientists score not much different than the average PhD but we do not.

The unique Terman study (1,500 >135 IQ children followed for decades) shows that the most / least successful had approximately the same IQ and scored the same in tests later in life. I.e. getting a great education did not yield extra IQ points in later tests; some 'Termites' accessed higher education *because of* their IQs. The differences in success were due to differences in personality (especially conscientiousness), not IQ. Cognitive ability is necessary for success but is not sufficient. Beyond a certain range, IQ *does* bring greater scientific success (publications, patents, breakthroughs) but *not* much more money. (Cf. A recent PhD thesis: '*Personality, IQ, and Lifetime Earnings*', Gensowski, 2013.)
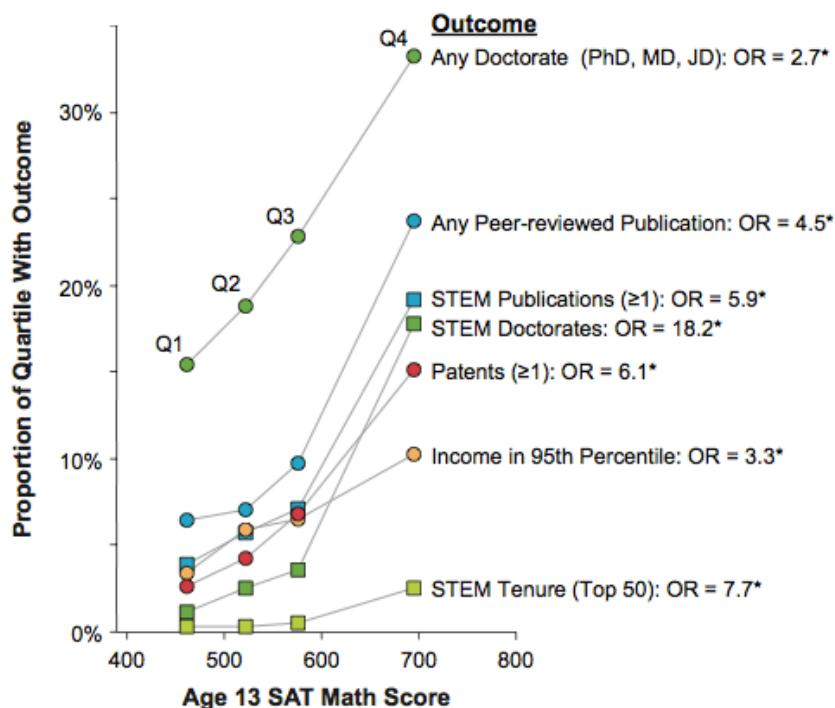
We can (imperfectly but much better than random) identify the 99+% of the population for whom significant success in science is very unlikely. Tests can predict future achievements for the extreme tail. Performance in the SAT-M test even at 13 predicts future scientific success and creativity, and can even distinguish within the top 1% performers; the top quartile of the 1% (1 in 10,000 children) are significantly more successful, so even a crude test at 13 (in which some score full marks so it is not fully testing the tail) has very useful predictive power. The fact that the relatively

---

[358] Given mathematicians and physicists regularly invade other fields with great success but the reverse almost never happens, 'it's pretty clear which things are easy or hard even for very smart people' (Hsu; cf. NYT on physicists invading genomics).

crude SAT-M test 'can identify 12 year olds who will earn this ultimate educational credential [PhD] at 50 times base rate is remarkable. Moreover, a 200 point difference in SAT scores by age 13 (500 versus 700) eventuates in marked differences by middle age in income, patents, refereed literary and scientific publications, and secured tenure track academic positions at top US universities…' (Lubinski). 'Using a quick IQ screen at age 17 probably allows you to identify future scientific leaders with 100 times better accuracy than chance, and future business leaders with 5-10 times better accuracy than chance. This follows immediately from (soft) 99+ and 90th percentile IQ thresholds. That is, the fraction of people in each of the two categories that is below the given threshold is very small.' (Hsu)

Re the most talented: 'adolescents scoring 500 or higher on SAT-M or SAT-V by age 13 (top 1 in 200), can assimilate a full high school course (e.g., chemistry, English, and mathematics) in three weeks at summer residential programs for intellectually precocious youth; yet, those scoring 700 or more (top 1 in 10,000), can assimilate at least twice this amount…' (Lubinski, 2010).

This chart shows the SMPY results (maths ability at 13) for the top 1% of maths ability broken down into quartiles 1-4: the top quartile of the top 1% clearly outperforms viz tenure, publication and patent rates.

Fig. 1. Accomplishments across individual differences within the top 1% of mathematical reasoning ability 25+ years after identification at age 13. Participants from Study of Mathematically Precocious Youth (SMPY) Cohorts 1, 2, and 3 (N = 2,385) are separated into quartiles based on their age-13 SAT-M score. The quartiles are plotted along the x-axis by their mean SAT-M score. The cutoff for a score in the top 1% of cognitive ability was 390, and the maximum possible score was 800. Odds ratios (OR) comparing the odds of each outcome in the *top* (Q4) and *bottom* (Q1) SAT-M quartiles are displayed at the end of every respective criterion line. An asterisk indicates that the odds of the outcome in Q4 was significantly greater than in Q1. STEM = science, technology, engineering, or mathematics. STEM Tenure (Top 50) = tenure in a STEM field at a U.S. university ranked in the top 50 by U.S. News and World Report's "America's Best Colleges 2007." Adapted in part from Park, Lubinski, and Benbow (2007, 2008).

http://infoproc.blogspot.com/2012/02/personnel-selection-horsepower-matters.html

Does pre-school work?

Most pre-school programmes, such as Head Start (USA) and Sure Start (UK), have shown no significant sustained improvement in IQ or test scores despite billions spent over years (or they

show brief improvements then 'fade out').[359] However, IQ is only part of what determines a child's future. As stated above re the unique Terman study of high IQ scientists, character, such as conscientiousness, is also vital. A small number of programmes that have been rigorously studied (particularly Perry,1962 and Abecedarian, 1972) have shown small improvements in test scores (which is what they were originally hoping to change) but, more importantly, long-term gains in employment, wages, health, criminality and so on. Studies of self-control clearly show that children with poor self-control are much more likely to be poor, have serious health problems (including dangerous addictions), and be criminals. Some programmes focused on 'executive functions', such as self-control, *have succeeded* and improved measures such as employment, crime, and drugs, though it has proved hard to scale up those few successes.[360]

Intelligence, wealth, and character

Most political discussion of education ignores both IQ and genetics. As Professor Plomin writes, 'educational research shies away from genetic explanations of individual variation [and] the basic pervasive level of genetics is strangely ignored'. Instead, political debates focus on class and wealth. It is assumed that wealthy people buy (directly or indirectly) better education and this translates into better exam scores including higher IQ scores? Is this true? Is IQ just a proxy for class? First, this argument cannot explain variation in IQ *within* families. Second, IQ predicts future educational achievement much better than class does. Third, shared environment does not have much impact on IQ and the impact *falls towards zero in adulthood (see above).*

'If we use a heritability range of .7-.9 (as deduced in twin studies; applicable for IQ in late adulthood and over a large range of childhood environments), then the resulting distribution of IQ differences for monozygotic twins would have SD = 5-8 points. *That is, clones raised in different families would typically have IQ differences of around 5-8 points.* (I.e. Pretty small considering a SD = 15 points.) N.B. very little of this difference is accounted for by SES - most of it is due to *non-shared* environmental effects. For a lower heritability value such as .5 (perhaps more appropriate for adolescents), the expected variation between clones would be about 10 points.' Hsu

> 'It is reasonable to consider the possibility that heritability of intelligence is higher in higher SES families because such families seem likely to provide more opportunities to realize differences in children's genetic potentials. Conversely, in lower SES families, genetic differences might be restrained by poverty… We attempted to replicate the finding that parental SES moderates the heritability of children's intelligence, with a greater genetic contribution to IQ in high-SES families compared to low-SES families. In a large UK-representative sample, we did not find evidence for the presence of such a gene-environment interaction across childhood and adolescence… Instead, using three different indices of SES, at eight ages from infancy through adolescence the emerging pattern appears to be one of *environment*-environment interaction rather than *gene*-environment interaction: *shared*

---

[359] Willingham underlines some recent studies claiming interventions that boost IQ: 1) 'infant formula supplemented with long chain polyunsaturated fatty acids boosts intelligence by about 3.5 points'; 2) 'interactive reading with a child raises IQ by about 6 points'; 3) 'preschools that include a specific language development component raise IQ scores by more than 7 points'; 4) intensive interventions with low-SES children to make the environment significantly more educationally enriching raised IQ scores by 4-7 points. Even if such results are replicated, *these gains are less than half a standard deviation* and would not make a difference between, for example, being able/unable to do a college degree.

[360] Cf. *Science (2011) special issue* on early years and Endnote on IQ. '*A gradient of childhood self-control predicts health, wealth, and public safety*' (Moffitt et al, PNAS, 2/2011). '*Preschools reduce early academic-achievement gaps: a longitudinal twin approach*' (Tucker-Drob, Psychological Science, 2/2012): a unique twin-study of pre-school effects. '*Why children succeed*' by Paul Tough (2012) summarises much of the literature including James Heckman's work. Is anybody doing a twin/adoption study involving children going to high quality pre-school *then* high quality schools, such as KIPP, to see whether the widely reported 'fade-out' of pre-school gains could be maintained by excellent schools?

experiences explain more of the variance in children's performance on IQ tests in more disadvantaged backgrounds. [Emphasis added.]

'Children from low-SES families face many physical and psychosocial environmental handicaps for their cognitive development. For example, low-SES children are read to less, have fewer books, less access to computers, and tend to watch more television. Parents tend to be less responsive to children in low-SES families, participate less in their children's school activities, and are more authoritarian. Children from more disadvantaged backgrounds tend to experience more instability, come from noisier, more crowded homes, and live in disadvantaged neighbourhoods with poorer facilities and inferior schools. To the extent that children growing up together experience these environments similarly, their cumulative effects are captured by [shared environment influences] in a twin model; experiences such as these seem likely to contribute to the observed greater variation in the cognitive ability performance of children from low-SES families.

'The notion that heritability may be lower in lower-SES families is appealing, in part because of its environmental implications: If heritability is lower in lower-SES families, it suggests that environmental interventions might be more effective in boosting cognitive development for children in lower-SES families. The present study, which is based on a large UK-representative sample of children followed longitudinally, leads to a similar implication. *Although the genetic influence on IQ is the same in lower-SES families, shared environmental influence appears to be greater in lower-SES families*, suggesting that family-based environmental interventions might be more effective in these families. However, two further aspects of the results temper the policy implications of this finding. *First, shared environmental influence is found in both lower- and higher-SES families and the difference in shared environmental influence between them is modest. Second, shared environmental influences on IQ decline from childhood to adulthood so that these influences might not have an impact in the long run.*' [Emphasis added] Plomin on class and IQ.

The finding that SES cannot account for IQ differences is connected to the Plomin and Daniels article (1987), one of - if not most - influential articles in the field of behavioural genetics, which showed that *once genetic relatedness is taken into account, siblings seem little more alike than strangers plucked randomly from the street and the variability in behaviour must come from the non-shared environment - not the shared environment of family, as almost everyone had assumed.*[361]

'In other words, despite a lifetime of proximity, your adopted child may bear no more similarity to you (in terms of, e.g., intelligence) than someone selected at random from the general population. The shared family environment that your children (biological or adopted) experience has little or no measurable effect on their cognitive development. While there *are* environmental effects on intelligence ([…] see here for Turkheimer's work suggesting low heritability in the case of severe deprivation), they seem to be idiosyncratic factors that can't be characterized using observable parameters such as the parents' SES, parenting style, level of education, or IQ. It is as if each child experiences their own random micro-environment, independent of these parental or family characteristics. The nonshared influences are by far

---

[361] This article is an update by Plomin of his famous paper about the almost untraceable effects of 'parenting style' and the 'nonshared environment'. The UK 'TEDS' long-term study of twins and schooling also touches on some of the issues concerning ability and schooling that most 'education debate' ignores.

the [largest] environmental (non-genetic) influences on intelligence - in fact, they are the *only* detectable non-genetic influences.' (Hsu)[362]

Analysis of differences in future education and earnings among siblings (ie. those in the same economic circumstances) shows that *ability has much more predictive power than class*. The data suggest that even 'utopian' plans to eliminate differences in wealth would not have big effects on future education and earnings. Analysis of data on gifted children shows that gifted children are healthier than privileged children *even though* the gifted are raised in homes >1SD *below* the privileged groups in terms of SES. Extensive analyses of US SAT scores shows that 'the vast majority of the test/academic performance relationship [i.e. that test scores predict future performance] was independent of SES'. We know that things like good teachers, parents reading to children, high status for education etc are correlated to academic performance, but these are also correlated with SES, and *controlling for SES has virtually no impact on the predictive power of test scores*.[363]

Gender differences… There are 'conspicuous gender differences in interest in people versus things' both in the general population and the SMPY cohort - about 1SD. 'Overall, boys and men are much more interested in working with things, gadgets, and inorganic material than girls and women are' (Lubinski). >70% of vet students and ~80% of psychology students are women. Among the SMPY cohort, men and women earn commensurate proportions of advanced degrees but men are more likely to take degrees in engineering and physical sciences, women are more likely to study law, medicine and social sciences. SMPY men are also more career focused, women prefer to spend less time at work. Lubinski argues that the choice of mathematically able women to apply their knowledge outside STEM subjects is arguably a good thing for society given the great need for mathematical analysis in other areas of human activity, and should not be automatically defined as a 'problem'.

Selection of elites… Cf. These blogs ([1],[2]) on a recent study on how firms select employees. There is an extreme focus on a super-elite of Harvard, Yale, Princeton, Stanford. Students from there are picked not because of the particular content of their education but because they are believed to have general cognitive abilities. However, lawyers, consultants, bankers, etc do *not* necessarily want the very brightest people from those schools because they are often considered 'nerds' who will not fit in socially. This blog, *Creators and Rulers*, discusses the difference between genuinely intelligent and talented 'creators' (scientists, tech start-ups) and the 'ruler' type that dominates politics and business organisations (CEOs with a history in law or sales) and seems very accurate to this author: 'the very existence of McKinsey, BCG, Bain, etc. is due to the mediocrity of the average Ruler. Think about what management consultants are selling and ask yourself why a senior exec. with 20 years experience in his industry needs a report prepared over 6 months by a 30 year old MBA.' Hsu distinguishes between 'hard elite' firms (hedge funds, VC, start-ups, technology) and 'soft elite' firms (law, consultancy, investment banks). 'Hard' performance is easier to measure (the

---

[362] 'I have no doubt that certain narrow skill sets (like piano playing, swimming, baseball) can be transmitted through parental effort. But the evidence seems to point the other way for intelligence, personality, religiosity and social attitudes. Are Chinese moms such outliers that they constitute a counterexample to the large-statistics studies cited above? It was shown by Turkheimer that in cases of extreme deprivation heritability can decrease significantly. Can Tiger Moms have the same effect at the positive end of the spectrum?' Hsu.

[363] Cf. These papers by [Lubinski] and [Sackett et al] on test scores and class. This 2012 [study] found that: 'a) the SAT retains the vast majority of its weight in predicting subsequent college grades when high school GPA and SES are controlled, b) the bivariate relationship between SAT score and freshman GPA is not substantially reduced when SES is controlled, and c) … it is not the case that the SAT is nothing more than a proxy for SES… [Further] the SES distribution of enrolled students at a college closely resembles the distribution of applicants to that college, which indicates that the barrier to college for low-SES students … is far more a matter of factors that affect the decision to enter the college application process than of exclusionary decisions on the part of colleges.'

gadget works/doesn't, but is the consultant's advice good?) and chimp politics matters less. 'Soft firms' recruit like they *know* what they do isn't that hard (they do not want the top scorer nerds) and outsource much of their decision to university recruiters who themselves perform less accurately than a simple weight-average algorithm.

Genes, class and social mobility… This paper on genes, class and social mobility in England 1800-2011 concludes: 1) regression to the mean suggests the Charles Murray dystopia of persistent upper/lower classes is *not* happening; 2) however, wealth persists longer than modern social studies suggest it should ('The 1800 underclass has already attained mediocrity. And the 1800 upper class will eventually dissolve into the mass of society, though perhaps not for another 300 years, or longer'); 3) the persistence of the wealthy born ~1800 as an educational elite is much stronger than the persistence of wealth, and ~¾ of educational status comes from characteristics inherited (biologically or culturally) from parents ('We may not be in the … Murray dystopia, but we are very close to it'); 4) 'the rate of regression to the mean for both wealth and educational status seems to have changed little over time, even though between 1800 and 2011 there have been enormous institutional changes in England… The modest effects of major institutional changes on social mobility implies that the important determination of persistence is transmission within families – either through genes or family environments. Indeed there almost seems to be a *social physics* here within families which controls the rate of regression to the mean, and makes it largely immutable from many outside institutional changes.' Hsu comments: 'It's possible that both Clark and Murray are correct. Clark's convergence may have stopped and reversed in the last 30 years due to globalization, meritocratic sorting, etc.' This paper answers some of the criticism of Clark's book. I do not know what the current scientific consensus is on Clark's arguments.

See Section 4 above for the BGI / Plomin project to identify the genes responsible for IQ.

'And what happens if one of these days people discover alleles for certain aspects of cognitive function? … Whatever ability you want, valued or not so valued, what if those alleles begin to come out? …What if somebody begins to look for the frequency of those alleles in different ethnic groups scattered across this planet? … the fact is if you look at the details of human evolution … you'll realize that most populations in humanity are the modern descendents of very small founder groups… So the fact is it's inescapable that different alleles are going to be present with different frequencies in different inbreeding populations of humanity or populations of humanity that traditionally have been genetically isolated from one another. It's not as if all the genes that we carry have been mixed with everybody else's genes freely over the last 100,000 years. Different groups have bred separately and have … founder affects and genetic drift, acquired different sets and different constellations of alleles.

'Then for the first time there could be a racism which is based not on some kind of virulent ideology, not based on some kind of kooky versions of genetics, because the eugenicists in the beginning of the 20th century, as well as the Nazis hadn't had any idea about genetics, they were just using the word, even though they knew nothing about the science of genetics as we understand it today. But what happens if now for the first time we, i.e., you who begin to understand genetics, begin to perceive that there are, in fact, different populations of humanity that are endowed with different constellation of alleles that we imagine are more or less desirable?

'What's going to happen then? I don't know. But some scientists say, well, the truth must come out and that everything that can be learned should be learned, and we will learn how to digest it and we will learn how

to live with that. But I'm not so sure that's the right thing. And you all have to wrestle with that as well.'
Weinberg, winner of the National Medal of Science 2007, Biology professor at MIT (2004)[364]

<center>*</center>

Genetics, school achievement, 'added value' measures

Twin studies on school performance have shown 'moderate heritability around 40-60%'. Studies of mathematical ability 'typically show high heritabilities around 60-70%'. Further, 'genetic influences appear to have largely generalist effects across diverse cognitive and academic abilities'. (Plomin's forthcoming *Genetic influence on educational achievement* examines the heritability of GCSE scores.)

'Added value' was introduced in 2002 on the basis that 'raw' attainment figures are 'largely an index of the calibre of pupil intake rather than any indication of the school's performance'.

Although it was intended to correct for natural ability, analysis of twin studies shows that 'about half of the variance of corrected-school achievement is due to genetic differences between children' - that is, even stripping out *g* or previous attainment does *not* strip out all genetic influence in general; *'achievement independent of ability may be just as heritable as achievement including ability because achievement is as much a function of genetically-driven appetites [such as motivation] as of aptitudes'.*

Therefore added value measures 'do not necessarily provide a pure indication of the school's effectiveness', 'show only minimal shared environmental contributions', and we should focus on investigating nonshared environments (i.e. not classrooms).

> 'We must conclude that it may not be possible to devise a measure of achievement that reflects only environmental influences [and] all aspects of achievement are suffused with genetic influence'.

> 'Instead of thinking about education as a way of countering genetic differences among children, the field of education might profit from accepting that children differ genetically in *how* and *how much* they learn. This way of thinking is compatible with the current trend towards personalizing education by optimizing children's learning, which is increasingly possible through the use of interactive information technology.'

> 'The opposite of personalized education is the attempt to use education to equalize children's learning, which if successful would have an unintended consequence: The resulting differences between children in their achievement would be even more heritable because one more source of environmental variation would be eliminated.'

Instead of thinking about education as *instruo* (build in) we should think of it as *educatio* (draw out).

---

[364] Socrates' nemesis, Callicles, rejected all conventional ideas about justice and virtue. 'Nature (*physis*) shows that … both among the other animals and in whole cities and races of men … justice (*dike*) [is] that the superior rule the inferior and have a greater share than they' (*Gorgias*). The majority 'like getting an equal share, since they are inferior' and they unsurprisingly wish to restrain the superior, those with intelligence (*phonesis*) and courage (*andreia*). Therefore the weak and inferior define justice as the restraint of greed (*pleonexia*) and use *nomos* (custom, law) to help them control the superior. *Nomos* is, therefore, a mere convention of the weak and inferior, which varies from city to city, and is less fundamental than *physis*. The superior man, he 'whose nature is equal to it', should trample the *nomos* of the inferior that violates nature's law and vent on the world his 'wantonness, lack of discipline, and freedom … [which] are excellence and happiness'.

Cf. '*Added Value Measures in Education Show Genetic as Well as Environmental Influence*', Haworth, Plomin *et al* (2011).

What are some of the effects of identifying the genetic basis of variance in cognitive abilities?

1) Progress identifying the genetic basis of variance in cognitive abilities, and therefore predicting phenotype from genotype, is the ultimate early warning system for common disorders such as reading difficulties and therefore could help identify who needs specific kinds of very early interventions, just as medicine hopes to prevent or limit problems emerging rather than treat them.

2) Further, when genes are found 'for' different kinds of ability that are 'polygenetic' (that is caused by lots of genes with small effects rather than a single mutation), the genes will be associated with the bottom, average, *and high* scorers; learning difficulties are quantitative dimensions rather than qualitative diagnoses. This finding could help change negative stereotypes about those with learning difficulties.

3) Because many genetic effects are general they blur distinctions between *ostensibly* different problems, such as reading and maths difficulties that have overlapping genetic causes; identifying 'generalist genes' will help with this. E.g. The fact that there are general maths genes means that a specific problem with one part of maths may indicate a problem with teaching rather than cognitive ability.

4) 'Because polygenetic liabilities are normally distributed as a bellshaped curve, they draw attention to the neglected positive side of the liability in addition to the negative side that has been the focus of most genetic research to date. The positive side of polygenic liability is not just low risk: It stimulates questions such as how children flourish rather than fail and about resilience rather than vulnerability' (Plomin, '*Child Development and Molecular Genetics*', 2013).

Endnote: A scientific approach to science teaching

Carl Wieman, Nobel winner for physics, has explored science teaching.

He thinks it is necessary, desirable and possible to make science education effective for a large fraction of the total population.

However, it is only possible if we approach the teaching of science like a science. There are three principles. 1. Practice must be based on objective data, including on how people learn. 2. Results must be disseminated in a scholarly manner and we must avoid 'reinventing the square wheel'; i.e. reinventing everything in a highly flawed form. 3. Apply technology properly.

Problems with traditional methods

Being a physicist, he decided to test how successfully he was teaching science. To his shock he discovered that *the vast majority of students did not understand what he was talking about*. His graduate students would arrive in his lab but after 17 years of success 'when they were given research projects to work on, *they were clueless* about how to proceed' and 'often it seemed they didn't even really understand what physics was.' However, after a few years in his lab they would be transformed into expert physicists. This was 'a consistent pattern' over many years.

Students do not remember what they hear in traditional lectures… A large number of studies shows that *students (including elite students at the best universities) remember little from traditional lectures.* This is unsurprising given 'one of the most well established — yet widely ignored — results of cognitive science: the extremely limited capacity of the short-term working memory. The research tells us that the human brain can hold a maximum of about seven different items in its short-term working memory and can process no more than about four ideas at once.'

Students do not master concepts taught in traditional lectures… There are good assessment tools. E.g. the Force Concepts Inventory test tests mastery of force and motion. Many courses now give this test at the start and end of the course to see how much has been learned. Analysis of results shows that students following a a traditional lecture course master 'no more than about 30 percent of the key concepts that they didn't already know at the start of the course' and this sub-30% finding is 'largely independent of lecturer quality, class size, and institution.' However, approaches 'involving more interactive engagement of students show consistently higher gains on the FCI and other tests.'

Students become more confused after a traditional course…You can interview people and place their beliefs on a spectrum between novice and expert. Novices see physics as isolated bits of information learned by memorisation if at all, and problem-solving is 'just matching the pattern of the problem to certain memorized recipes.' 'Experts — i.e., physicists — see physics as a coherent structure of concepts that describe nature and that have been established by experiment. Expert problem-solving involves employing systematic, concept-based, and widely applicable strategies.' One hopes that a physics course moves people at least a bit from novice towards expert. 'On average, students have more novicelike beliefs after they have completed an introductory physics course than they had when they started; this was found for nearly every introductory course measured. More recently, my group started looking at beliefs about chemistry. If anything, the effect of taking an introductory college chemistry course is even worse than for taking physics.'

So: 'traditional science instruction … is explicitly built around teaching concepts and is being provided by instructors who, at least at the college level, are unquestionably experts in the subject. And yet their students are not learning concepts, and they are acquiring novice beliefs about the subject.' How? Cognitive science shows that as well as 'lots of factual knowledge', 'experts have a mental organizational structure that facilitates the retrieval and effective application of their knowledge' and 'an ability to monitor their own thinking ("metacognition")…They are able to ask themselves, "Do I understand this? How can I check my understanding?"'

Traditional teaching focuses on facts and assumes that expert-like ways of thinking 'come along for free' but 'this is not what cognitive science tells us.' Instead, 'students need to develop these different ways of thinking by means of extended, focused mental effort.' (Understanding science resides in long-term memory which is developed 'via the construction of assembly of component proteins' so mastery is impossible without an extended mental construction process.)

Why did Wieman's graduate students become experts? 'A lot of educational and cognitive research can be reduced to this basic principle: People learn by creating their own understanding. But that does not mean they must or even can do it without assistance. Effective teaching facilitates that creation by getting students engaged in thinking deeply about the subject at an appropriate level and then monitoring that thinking and guiding it to be more expert-like.' This is what his graduate students are doing. 'They are focused intently on solving real physics problems, and I regularly probe how they're thinking and give them guidance to make it more expert-like. After a few years in that environment they turn into experts, not because there is something magic in the air in the

research lab but because they are engaged in exactly the cognitive processes that are required for developing expert competence.'

How can we operationalise this knowledge?

Unfortunately, we cannot put everyone into a lab to work closely with a faculty member [though this is roughly what the famous Russian maths and physics schools did and why they succeeded in preparing pupils for university research].

First, reduce cognitive load. 'Anything the teacher can do to reduce that cognitive load while presenting the material will help. Some ways to do so are obvious, such as slowing down. Others include having a clear, logical, explicit organization to the class (including making connections between different ideas presented and connections to things the students already know), using figures where appropriate rather than relying only on verbal descriptions and minimizing the use of technical jargon.'

Second, explain why the topic is studied, how it operates in the real world, why it makes sense, and how it connects to things already known. Wieman has tested doing this and it works.

Third, teachers must engage students, monitor their thinking, and give feedback. 'To do this effectively, teachers must first know where the students are starting from in their thinking, so they can build on that foundation. Then they must find activities that ensure that the students actively think about and process the important ideas of the discipline. Finally, instructors must have mechanisms by which they can probe and then guide that thinking. This takes much more than just mastery of the topic — it requires, in the memorable words of Lee Shulman, "pedagogical content knowledge."… This classroom experience has to be followed up with extended "effortful study," where the student spends considerably more time than is possible in the classroom developing expert-like thinking and skills.'

'Extended, highly focused mental processing is required to build those little proteins that make up the long-term memory. No matter what happens in the relatively brief period students spend in the classroom, there is not enough time to develop the long-term memory structures required for subject mastery. To ensure that the necessary extended effort is made, and that it is productive, requires carefully designed homework assignments, grading policies, and feedback. As a practical matter, in a university environment with large classes the most effective way for students to get the feedback that will make their study time more productive and develop their metacognitive skills is through peer collaboration.'

Use of technology

1. Education purposes must drive technology use, not vice versa. 2. Technology use must be tested.

Some proven ideas…

A. '"Just-in-time teaching" … uses the Web to ask students questions concerning the material to be covered, questions that they must answer just before class. The students thus start the class already engaged, and the instructor, who has looked at the students' answers, already knows a reasonable amount about their difficulties with the topic to be covered.'

B. 'Clickers'… 'Each student has a clicker with which to answer questions posed during class. A computer records each student's answer and can display a histogram of those responses. The clicker efficiently and quickly gets an answer from each student for which that student is

accountable but which is anonymous to their peers. I have found that these clickers can have a profound impact on the educational experience of students. The most productive use of clickers in my experience is to enhance the Peer Instruction (PI) technique developed by Eric Mazur, particularly for less active or assertive students. [See below]

'I assign students to groups the first day of class (typically three to four students in adjacent seats) and design each lecture around a series of seven to 10 clicker questions that cover the key learning goals for that day. The groups are told they must come to a consensus answer (entered with their clickers) and be prepared to offer reasons for their choice. It is in these peer discussions that most students do the primary processing of the new ideas and problem-solving approaches. The process of critiquing each other's ideas in order to arrive at a consensus also enormously improves both their ability to carry on scientific discourse and to test their own understanding.

'Clickers also give valuable (albeit often painful) feedback to the instructor when they reveal, for example, that only 10 percent of the students understood what was just explained. But they also provide feedback in less obvious ways. By circulating through the classroom and listening in on the consensus-group discussions, I quickly learn which aspects of the topic confuse students and can then target those points in the follow-up discussion. Perhaps even more important is the feedback provided to the students through the histograms and their own discussions. They become much more invested in their own learning. When using clickers and consensus groups, I have dramatically more substantive questions per class period — more students ask questions and the students represent a much broader distribution by ethnicity and gender — than when using the peer-instruction approach without clickers.'

C. Online interactive simulation. 'This technique can be highly effective and takes less time to incorporate into instruction than more traditional materials. My group has created and tested over 60 such simulations and made them available for free (www.phet.colorado.edu). We have explored their use in lecture and homework problems and as replacements for, or enhancements of, laboratories.

'The "circuit construction kit" is a typical example of a simulation. It allows one to build arbitrary circuits involving realistic-looking resistors, light bulbs (which light up), wires, batteries, and switches and get a correct rendition of voltages and currents. There are realistic volt and ammeters to measure circuit parameters. The simulation also shows cartoonlike electrons moving around the circuit in appropriate paths, with velocities proportional to current. We've found this simulation to be a dramatic help to students in understanding the basic concepts of electric current and voltage, when substituted for an equivalent lab with real components.

'As with all good educational technology, the effectiveness of good simulations comes from the fact that their design is governed by research on how people learn, and the simulations are carefully tested to ensure they achieve the desired learning. They can enhance the ability of a good instructor to portray how experts think when they see a real-life situation and provide an environment in which a student can learn by observing and exploring. The power of a simulation is that these explorations can be carefully constrained, and what the student sees can be suitably enhanced to facilitate the desired learning.'

Results

Table 1. Comparison of Learning Results from Traditionally Taught Courses and Courses Using Research-Based Pedagogy

Traditional Instruction	Research-Based Instruction
Retention of information from lecture: 10% after 15 minutes	Retention of information from lecture: More than 90 % after 2 days
Gain in conceptual understanding: 25%	Gain in conceptual understanding: 50-70%
Beliefs about physics and problem-solving: significant drop	A small improvement

What to do?

It is necessary to change college practice towards a scientific approach because colleges teach the next generation of science teachers. There are two big challenges.

In universities 'there is generally no connection between the incentives in the system and student learning... The real problem is that we have almost no authentic assessments of what students actually learn, so it is impossible to broadly measure that learning and hence impossible to connect it to resources and incentives. We do have student evaluations of instructors, but these are primarily popularity contests and not measures of learning.'

'The second challenge is that while we know how to develop the necessary tools for assessing student learning in a practical, widespread way at the university level, carrying this out would require a significant investment. Introducing effective research-based teaching in all college science courses—by, for instance, developing and testing pedagogically effective materials, supporting technology, and providing for faculty development—would also require resources. But the budget for R&D and the implementation of improved educational methods at most universities is essentially zero. More generally, there is not the political will on campus to take the steps required to bring about cultural change in organizations like science departments.'

Cf. http://www.cwsei.ubc.ca/resources/files/Wieman-Change_Sept-Oct_2007.pdf

Eric Mazur (Harvard), like Wieman, thought he was a success until he came across tests of what physics undergraduates understand and gave the tests to his own students. These showed that although he was teaching students successfully some mathematical and problem-solving techniques, they did not finish the course understanding physics but reverted to Aristotelian logic.

He stumbled across the idea of *peer instruction*. 'First, when one student has the right answer and the other doesn't, the first one is more likely to convince the second - it's hard to talk someone into the wrong answer when they have the right one. More important, a fellow student is *more likely* to reach them than Professor Mazur - and this is the crux of the method. You're a student and you've only recently learned this, so you still know where you got hung up, because it's not that long ago that *you* were hung up on that very same thing. Whereas Professor Mazur got hung up on this point when he was 17, and he no longer remembers how difficult it was back then. He has lost the ability to understand what a beginning learner faces.'

'Interactive learning triples students' gains in knowledge as measured by the kinds of conceptual tests that had once deflated Mazur's spirits, and by many other assessments as well. It has other salutary effects, like erasing the gender gap between male and female undergraduates… There's also better retention of knowledge. "In a traditional physics course, two months after taking the final exam, people are back to where they were before taking the course," Mazur notes. "It's shocking." (Concentrators are an exception to this, as subsequent courses reinforce their knowledge base.) Peer-instructed students who've actively argued for and explained their understanding of scientific concepts hold onto their knowledge longer. Another benefit is cultivating more scientists. A comparison of intended and actual concentrators in STEM … fields indicates that those taught interactively are only half as likely to change to a non-STEM discipline as students in traditional courses.'

'Scores of Harvard faculty members are experimenting with innovative styles of teaching in their classes. Mazur's profile is perhaps the highest because he has been at it for two decades and has poured so much passion into the pursuit. But across the University's faculties, instructors are trying out new, more effective modes of pedagogy.'

'Rigorous evaluations using statistical analysis can help distinguish the most promising innovations. For his part, Mazur has collected reams of data on his students' results. (He says most scholars, even scientists, rely on anecdotal evidence instead.) End-of-semester course evaluations he dismisses as nothing more than "popularity contests" that ought to be abolished. "There is zero correlation between course evaluations and the amount learned," he says. "Award-winning teachers with the highest evaluations can produce the same results as teachers who are getting fired."

'"More and more faculty are finding that the traditional lecture no longer suits them," Aladjem notes. "And they are finding alternative ways to connect with students. Some are quite sophisticated in using course websites, blogs, and other means to be in touch. Michael McCormick, Goelet professor of medieval history, holds office hours late at night via Skype, and it's very popular. Nearly all undergraduates have laptops, smart phones, or other computing devices and use them all constantly…

'"We want to educate leaders, the innovators of society," Mazur says. "Let's turn our students into *real* problem solvers. In a real-world problem, you know where you want to get, but you don't know how to get there. For example: how can I bake a cake with no flour? The goal is known, but the prescription to get there isn't. Most tests and exams at Harvard are not like that; they are questions where you need to determine what the *answer* is. In physics it might be, What was the velocity of the car before it hit the tree? There, you know exactly what you need to do: you *have* a prescription to calculate velocity, but you don't know the velocity. It's the opposite of a real-life problem, because you know the prescription, but you don't know the answer.

'"Even now, if I give my students a problem on an exam that they have not seen before, there will be complaints: 'We've never done a problem of this kind.' I tell them, 'If you *had* done a problem of this kind, then by definition, this would not be a problem.' We have to train people to tackle situations they have not encountered before. Most instructors avoid this like the plague, because the students dislike it. Even at Harvard, we tend to keep students in their comfort zone. The first step in developing those skills is stepping into unknown territory.

'"It's not easy. You get a lot of student resistance," he continues. "You should see some of the vitriolic e-mails I get. The generic complaint is that they have to do all the learning *themselves.*

Rather than lecturing, I'm making them prepare themselves for class — and in class, rather than telling them things, I'm asking them questions. They'd much rather sit there and listen and take notes. Some will say, 'I didn't pay $47,000 to learn it all from the textbook. I think you should go over the material from the book, point by point, in class.' Not realizing that they learn precious little by that, and they should actually be offended if I did that, because it's an insult to their intelligence — then, I'm essentially reading the book to them.'"

http://harvardmagazine.com/2012/03/twilight-of-the-lecture

http://bokcenter.harvard.edu/icb/icb.do?keyword=k1985&tabgroupid=icb.tabgroup122049

<center>*</center>

'*Retrieval Practice Produces More Learning than Elaborative Studying with Concept Mapping*', by Jeffrey D. Karpicke and Janell R. Blunt (Science, Jan 2011)

'Most thought on human learning is guided by a few tacit assumptions. One assumption is that learning happens primarily when people encode knowledge and experiences. A related assumption is that retrieval - the active, cue-driven process of reconstructing knowledge - only measures the products of a prior learning experience but does not itself produce learning. Just as we assume that the act of measuring a physical object would not change the size, shape, or weight of the object, so too people often assume that the act of measuring memory does not change memory. Thus most educational research and practice has focused on enhancing the processing that occurs when students encode knowledge – that is, getting knowledge "in memory". Far less attention has been paid to the potential importance of retrieval to the process of learning.

'… However, research in cognitive science has challenged the assumption that retrieval is neutral and uninfluential in the learning process. Not only does retrieval produce learning, but a retrieval event may actually represent a more powerful learning activity than an encoding event.

'This research suggests a conceptualization of mind and learning that is different from one in which encoding places knowledge in memory and retrieval simply accesses that stored knowledge. Because each act of retrieval changes memory, the act of reconstructing knowledge must be considered essential to the process of learning...

'Most prior research has not used assessments thought to measure meaningful learning, which refers to students' abilities to make inferences and exhibit deep understanding of concepts. But perhaps the greatest impediment to broad application of retrieval practice is that we do not know whether retrieval activities are more effective than other active, elaborative learning activities. Retrieval practice might produce levels of learning that are essentially the same as those produced by elaborative studying. Alternatively, if there are retrieval-specific mechanisms that promote learning, then retrieval practice may represent a way to promote student learning that goes beyond elaborative study activities used in STEM education...

'We examined the effectiveness of retrieval practice relative to elaborative studying with concept mapping. In concept mapping, students construct a diagram in which nodes are used to represent concepts and links connecting the nodes represent relations among the concepts. Concept mapping is considered an active learning task, and it serves as an elaborative study activity when students construct concept maps in the presence of the materials they are learning...

[Experiments tested undergraduates using 1) elaborative studying plus concept mapping and 2) studying and retrieving. (2) performed much better as a method when tested using tests a week later.]

'[There was] about a 50% improvement in long-term retention scores… [Other tests were done confirming the superiority of retrieval.]

'Students predicted that repeated studying would produce the best long-term retention and that practicing retrieval would produce the worst retention, even though the opposite was true…

'Retrieval is not merely a read out of the knowledge stored in one's mind – the act of reconstructing knowledge itself enhances learning. This dynamic perspective on the human mind can pave the way for the design of new educational activities based on consideration of retrieval processes.'

*

Endnote: Gigerenzer - Bayes, risk, and education

'My ventures are not in one bottom trusted
Nor to one place; nor is my whole estate
Upon the fortune of this present year;
Therefore, my merchandise makes me not sad.'
Antonio, Merchant of Venice, 1,i.

Bayes' Theory and 'natural frequencies'…

'Problem: Given that the number of times in which an unknown event has happened and failed: Required the chance that the probability of its happening in a single trial lies somewhere between any two degrees of probability that can be named.' Thomas Bayes's *'Essay Towards Solving A Problem In The Doctrine of Chances'* was not published until 1764, after his death, and remained in obscurity for another 20 years.

Many problems of conditional probability ('what is the probability of X given Y?' or $P(X|Y)$) are presented in terms of Bayes' Theory using percentages: $P(A|B) = [P(B|A) \times P(A)] / P(B)$.

E.g. *'The probability that a woman of age 40 has breast cancer is about 1%. If she has breast cancer, the probability that she tests positive on a screening mammogram is 90%. If she does not have breast cancer, the probability that she nevertheless tests positive is 9%. What are the chances that a woman who tests positive actually has breast cancer?'*

The answer is ~9%, or 1 in 111.[365] However, many studies show that there is a huge variability in answers (from ~1%-90%) and the median estimate of *doctors* (never mind patients) on the problem above is ~70%. This is frighteningly far from the truth.

Another way to present the same problem is:

'Think of 100 women. One has breast cancer, and she will probably test positive. Of the 99 who do not have breast cancer, 9 will also test positive. Thus, a total of 10 women will test positive. How many of those who test positive actually have breast cancer?'

---

[365] Using Bayes: P(Disease|Positive test) = P(P|D)xP(D)/P(P) = (0.9x0.01) / (0.01x0.9) + (0.99x0.09) = 0.0917… ie. ~9%.

Here it is much simpler to see the answer is ~10%. When the same problem is presented as here in *natural frequencies*, rather than as percentages to be plugged into Bayes' equation, most doctors are close to the right answer and the huge variability largely disappears.

Tests show that computer training systems to teach children how to turn probabilities into natural frequencies are much more effective than traditional methods and, even more dramatically, the lessons are remembered longer. It even seems that students really internalise the new methods and use them because they show little decay in performance over time. Computer training is also popular with the students and can reverse hostility to maths learning.

Three schools of probability…

There are three broad schools of probability. 1) *'Degrees of belief'* (or 'subjective probabilities'). Here one can apply a probability to almost anything provided one follows basic formal rules (like the probability of exhaustive and exclusive events sums to 1). E.g. The doctor who did the first heart transplant put the odds of success at '80 percent' (the patient died after 18 days). 2) *'Propensities'*. E.g. The design of a fair dice means each number has a 1/6 chance of showing. If a design is unknown you cannot assign a numerical probability. 3) *'Frequencies'*. For a frequentist, probability is the relative frequency of an event in a specified relative class based on a number of trials. E.g. DASA describes the risk of failure of an Ariane rocket as 0.4% based on its conclusions about its parts (a 'propensity' view) but a frequentist would say it is 8/94 based on the actual number of failures. Knight's famous distinction between 'uncertainty' and 'risk' was actually a distinction between 'subjective probabilities' and 'frequentism'. (Gigerenzer).

The problem of 'relative risk'…

One often reads media accounts such as 'drug X reduces the risk of death by 22%'. What does this mean? Many think it means something like 'of 1,000 people 220 will avoid death'. In fact, it is usually based on something like 'of the group who took the drug 32 died while of the control group 41 died' (9/41=22%). In the same way that it is better to turn percentages into natural frequencies to improve understanding, it is better to turn medical reports into '*number needed to treat*' to save one life (NTT): the NTT of drug X is ~111 (9/1000 $\approx$ 1/111 so 111 need to take the drug to save 1 person). In Britain recently, an increase in the risk of thromboembolism (blood clot) due to the pill was presented in relative risk terms of a 'doubling' of risk, which actually reflected an increase from 1/14,000 to 2/14,000. The consequence was huge anxiety and a rise in the number of unwanted pregnancies and abortions. A study of Anglia and Oxford regional health authorities showed that people responsible for purchasing decisions were unable to realise that the *same* underlying information was presented in four different ways, and they were strongly influenced by the 'relative risk' representation. Only 3/140 realised the four representations reflected the same results.

Learning from disaster…

Medical errors have proved harder to limit than in airlines. Incentives for doctors and others is not the same as for pilots; they do not die from their own errors. Although evidence-based medicine is spreading, it was a long slog. There was huge resistance to randomised control trials (pioneered by Fisher in the 1920s). Many doctors did not let patients see their records. Katz's '*The Silent World of Doctor and Patient*' (1984) argued that patients should be involved more in decisions affecting them and was almost uniformly attacked by surgeons. Many doctors do not think telling patients the truth is wise; many think that a chat and prescription triggers the placebo effect so is worthwhile even if the patient is clueless about what has been said; some complain it can have perverse effects like encouraging them to visit faith healers. Most patients do not understand the statistics which

are often presented confusingly. Few understand the dangers of screening and many think it operates like a vaccine. Doctors fear that inaction risks legal action and are incentivised by the fact that action is paid; they fear the emotions of missing cancer more than the emotions of advising action that proves unnecessary.

There are similar problems in law. E.g. Dershowitz got away with very misleading statistics about conditional probability in the OJ Simpson case. Both sides in the OJ case focused on the wrong probability - (P(man murdered wife | she's battered), which is ~1/2,500) instead of the right one: P(man murdered wife | she's battered *and is murdered*), which is ~9/10 (Gigerenzer and Strogatz p. 187-8). The accuracy of DNA is routinely overstated. ⅔ of predictions regarding the likelihood of prisoners committing violence are wrong. Only one out of 175 US law schools has a course in basic statistics.

Here is a summary of some of the arguments between Bayesians and 'frequentists': '*A 250-year argument*', Efron, AMS 2013. Gigerenzer's *Reckoning With Risk* is a great introduction to Bayes.

Endnote: Cognitive Biases and Decisions

Conformity and Groupthink…

Three classic experiments have shown the powerful evolved instinct to *conform* to a group view, a flip-side of our evolved in-group *solidarity*.

1) Asch showed how people will force themselves to conform to group ideas about the length of lines on a page despite the group view being obviously wrong (1951, repeated recently with fMRI scanners).

2) Milgram showed that people will persist in giving what they think are fatal electric shocks to people, with actors' screams ringing in their ears, on the grounds that it is their 'duty': 'ordinary people, simply doing their jobs, and without any particular hostility on their part, can become agents in a terrible destructive process. Moreover, even when the destructive effects of their work become patently clear, and they are asked to carry out actions incompatible with fundamental standards of morality, relatively few people have the resources needed to resist authority… Morality does not disappear - it acquires a radically different focus: the subordinate person feels shame or pride depending on how adequately he has performed the actions called for by authority' Milgram (1963).

3) Zimbardo created a fake prison with subjects assigned to the roles of prisoners and guards, leading rapidly to extreme abuse (1971).

The optimistic finding is that even lone dissenters can change grotesque dynamics radically; the pessimistic finding is that few of us can cope with the emotional trauma of dissent. 'The bearers of bad news tend to become pariahs, shunned and ignored by other employees. When pessimistic opinions are suppressed, while optimistic ones are rewarded, an organization's ability to think critically is undermined. The optimistic biases of individual employees become mutually reinforcing, and unrealistic views of the future are validated by the group.' Kahneman.[366]

The Condorcet Jury Theorem…

---

[366] A recent paper, '*Contesting the nature of conformity*' (Haslam & Reicher, 2012), argues that the Milgram / Zimbardo experiments have been consistently misinterpreted and the 'banality of evil' hypothesis is misleading.

Suppose a group of people are answering the same question either true/false (T/F) and the probability that each voter will be T exceeds 50%. Then the probability of a correct answer by the majority of the group increases toward 100% as the size of the group grows. (However, if each is less than 50% likely to be T, then the probability of the group answer being T falls toward zero as the group grows.) This depends on two conditions: 1) each person is more than 50% likely to be right; 2) majority rule is used. Condorcet assumed: 1) people would be unaffected by whether their vote is decisive; 2) people would be unaffected by others' votes; 3) the probability that a group member would be right is statistically unrelated to whether another is right. The first two hold for 'statistical groups'. The third may be violated if people see things the same way. Further, a group can get the right answer even if not everybody is >50% likely to be right. Even if only 40% of the group are >50% likely to be right and 60% are 50-50 *with the errors random*, then the group, if large enough, will still produce a plurality for T (cf. *Who Wants To Be A Millionaire?*). If errors are worse than random, perhaps because of a systematic bias, then the statistical group answer will be wrong. Condorcet warned that 'the prejudices to which this voter is subject' mean 'it can be dangerous to give a democratic constitution to an unenlightened people.'[367] (*Infotopia*, Sunstein)

Some questions will get the most accurate answer from a statistical group but many questions are impenetrable to most people, even very well-educated people (e.g. Sunstein's survey of the Chicago Law School faculty on the question of the weight of the space shuttle's fuel produced a hopeless median and mean). Averages of groups of experts will produce a better result in many fields. From Aristotle to Habermas and Rawls ('Discussion is a way of combining information and enlarging the range of arguments'), academics have been optimistic that large groups will produce good reasoning but it depends on the question and the rules of the voting. Many studies have shown that deliberating groups can be very poor at aggregating information, can suppress the diversity of views, can overvalue information held in common, and increase confidence in a wrong decision. In some cases they perform as well as or slightly better than the average of individual views but worse than the best individuals. Groups do well with crosswords *because it is easy to spot right and wrong answers.* Anything that touches on group identity and cohesion potentially distorts discussion in extreme ways.

Endnote: Some sketches of the history of economic thinking[368]

In 1987, the Santa Fe Institute, founded by Gell Mann and others, organised a ten day meeting to discuss economics. On one side, they invited leading economists such as Kenneth Arrow and Larry Summers; on the other side, they invited physicists, biologists, and computer scientists, such as Nobel-winning Philip Anderson and John Holland (inventor of genetic algorithms). When the economists explained their assumptions, Phil Anderson said to them, '*You guys really believe that?*' One physicist later described the meeting as like visiting Cuba - the cars are all from the 1950's so on one hand you admire them for keeping them going, but on the other hand they are old technology; similarly the economists were ingeniously using 19th Century maths and physics on very out-of-date models. The physicists were shocked at how the economists were content with simplifying assumptions that were obviously contradicted by reality, and they were surprised at the way the economists seemed unconcerned about how poor their predictions of reality. This note summarises some of the history.

1. The Marginalist Revolution

---

[367] 'A very numerous assembly cannot be composed of very enlightened men. It is even probable that those comprising this assembly will on many matters combine great ignorance with great prejudices. Thus there will be a great number of questions on which the probability of the truth of each voter will be below ½. It follows that the more numerous the assembly, the more it will be exposed to the risk of making false decisions' (Condorcet).

[368] Much of this section comes from Beinhocker (2005), Mandelbrot (2005), and Buchanan (2013).

2. Von Neumann, 'equilibrium', and the invention of Game Theory

3. General Equilibrium Theory and the 'Neoclassical synthesis'

4. Orthodox Financial Theory

5. Some conclusions

<center>*</center>

1. The Marginalist Revolution (late 19th Century)…

Bernoulli and 'utility'… In 1738, Daniel Bernoulli (nephew of Jacob Bernoulli who first introduced the Law of Large Numbers) had published a revolutionary paper that introduced the concept of 'utility'.

> 'Ever since mathematicians first began to study the measurement of risk, there has been general agreement on the following proposition: *Expected values are computed by multiplying each possible gain by the number of ways in which it can occur, and then dividing the sum of these products by the total number of cases.*'

However, price and probability are insufficient for judging an unknown future. While the objective facts concerning expected value are the same for everyone, utility varies yet it can be measured:

> '... the *value* of an item must not be based on its *price* but rather on the *utility* that it yields... [Utility] is dependent on the particular circumstances of the person making the estimate...

> '[T]he utility resulting from any small increase in wealth will be inversely proportional to the quantity of goods previously possessed... Considering the nature of man, it seems to me that the foregoing hypothesis is apt to be valid for many people to whom this sort of comparison can be applied.'

Bentham (1748-1832) rediscovered Bernoulli's idea of utility at the end of the 18th Century.[369] 'Classical' economics developed a theory of how utility and price are connected in the economy. While Turgot had argued that there are *diminishing marginal returns on production*, Gossen argued that there is a *law of diminishing marginal utility on consumption* - eventually one will not keep wanting to add another X to one's pile and the value of another X falls. Price therefore finds a balance between diminishing returns and diminishing utility, between production and consumption, and between supply and demand.

From Newton to 'equilibrium' theory...

---

[369] In *The Principles of Morals and Legislation*, he wrote: 'Nature has placed mankind under the governance of two sovereign masters, *pain* and *pleasure*. It is for them alone to point out what we ought to do, as well as to determine what we shall do... The *principle of utility* recognises this subjection, and assumes it for the foundation of that system, the object of which is to rear the fabric of felicity by the hands of reason and law... [Utility is] that property in any object, whereby it tends to produce benefit, advantage, pleasure, good, or happiness ... when the tendency it has to augment the happiness of the community is greater than any it has to diminish it.' Von Neumann said of Bentham's famous idea that we should pursue 'the greatest possible good for the greatest possible number': '[it is] taken literally, self-contradictory... A guiding principle cannot be formulated by the requirement of maximising two (or more) functions at once.'

In the second half of the 19th Century, the extraordinary development of physics since Newton spawned 'positivism' - a movement to incorporate the methodology of physics into various intellectual areas, including economics. While Marx developed what he claimed was a 'scientific' economics, others applied the methods of physics directly to create new mathematical models of the economy. The birth of Neoclassical economics is conventionally dated to the work of Walras and Jevons in the 1870's in what is sometimes called 'the marginal utility revolution'.

Newtonian mechanics had shown that in the absence of forces acting on a system, it will remain in *equilibrium*. In the 19th Century, Lagrange showed that the laws of Newtonian physics implied a principle of 'least action'. Cournot believed that a similar principle of equilibrium could be applied to the laws of supply and demand - *there must be a point at which all the goods traded in a market would find their equilibrium.*

Leon Walras had failed to get into the *Ecole Polytechnique* because of insufficiently good maths; he also failed as an engineer and a novelist. Influenced by a walk in the garden with his father, a French book on physics (*Elements of Statics* by Poinsot), and by Cournot, he decided to make economics a 'science' in the spirit of classical physics and build a mathematical model to show how equilibrium could occur. In 1872, he finished his *Elements of a Pure Economics*.

Walras wanted to be able to predict economic states as physicists could predict various physical states. He therefore imported the use of *calculus* and the concept of *equilibrium*. Equilibrium occurs when the forces acting on an object cancel each other out. A *stable equilibrium* is something like when a marble is dropped into a bowl, wobbles around then settles in the basin where it will stay unless some external force disturbs it. An *unstable equilibrium* is like a pencil balanced on its end; it is easily broken. Some systems have *multiple equilibria*. To build a mathematical model, Walras took from Poinsot's book the maths in the chapter 'On conditions of equilibrium expressed by means of equations'.

Walras' model is based on the assumption that when the economy is in equilibrium, nobody wants to trade because their utility is satisfied by the current distribution of goods; when people are satisfied, they stop trading and there is what he called 'the general equilibrium point'. The question therefore becomes how to find the price levels which produce a state in which everyone has what they want. He assumed an *auctioneer* and that one good is used as a *pseudo-currency*; the auctioneer would then run through the list of goods taking bids from the floor and settling on a price for each good; at the end, all the prices would be set and everyone trades at the agreed prices and the *general equilibrium point* would be reached (which Walras called *tâtonnement*, groping). This model was based on rational pursuit of self-interested utilities. Walras argued that his 'pure theory of economics is a science which resembles the physico-mathematical sciences in every respect' (Beinhocker pp. 29-33).

> '[A]ll these results are marvels of the simple application of the language of mathematics to the quantitative notion of need or utility… You can be sure that the economic laws that result from it are just as rational, just as precise and just as incontrovertible as were the laws of astronomy at the end of the 17th Century.' Walras.

In 1867 Sir William Thomson (Lord Kelvin) and Peter Tait published *A Treatise on Natural Philosophy* describing recent developments in physics concerning Maxwell and 'fields' such as electromagnetism. *William Stanley Jevons*, a scientist and mathematician who for financial reasons worked in finance instead of science, read it and applied aspects of it to economics. His *Theory of Political Economy* (1871) developed a mathematical model of the economy (which he called 'a calculus of moral effects') using the equations of field theory in which, just as a pendulum's motion

is determined by the force (gravity) and the constraints upon it (length of string), so an economy is determined by a force, in the form of self-interest driving the pursuit of utility ('an attraction between a wanting being and what is wanted'), and the constraints upon it, in the form of finite resources.

> 'The notion of value is to our science what that of energy is to mechanics....
>
> '... value depends entirely upon utility ...We have only to trace out carefully the natural law of the variation of utility, as depending upon the quantity of a commodity in our possession, in order to arrive at a satisfactory theory of exchange... Previous to the time of Pascal, who would have thought of measuring doubt and belief?... Now there can be no doubt that pleasure, pain, labour, utility, value, wealth, money, capital etc are all notions admitting of quantity; nay, the whole of our actions in industry and trade certainly depend upon comparing quantities of advantage and disadvantage.'

Beinhocker writes:

> 'Jevons' lasting contribution was to portray the problem of economic choice as an exercise in constrained optimization. That is, given the amounts available, a consumer will calculate what quantities of various goods will make him or her the most happy… [D]ifferences in individual utilities create a kind of potential energy for trade.'

The theory of 'marginal utility' posits that people make decisions based on how they subjectively value units of something 'at the margin'. A famous analogy explains it thus: if you have five sacks of grain, no chances for trade, and various uses for it ranging from feeding yourself (most important) to feeding parrots to keep you entertained (least important), then if you lose one sack you will not reduce all uses equally but will let the parrots starve; the parrots are 'on the margin' because their relative utility is lowest. The marginal utility of an object is therefore its least important role. Marginal cost is the amount of money needed to increase production by one more unit.

Beinhocker recounts how in 1901 Walras sent the legendary mathematician Poincaré a copy of *Elements of a Pure Economics.* Poincaré replied that he was not hostile 'a priori' to the use of mathematics in economics 'as long as one does not go beyond certain limits' and warned that Walras' assumptions could render the theory 'devoid of all interest'.

> '[Y]ou regard men as infinitely selfish and infinitely farsighted. The first hypothesis may perhaps be admitted in a first approximation, the second may call for some reservations.' (Poincaré to Walras)

Vilfredo Pareto defined four kinds of trades: both gain, win-win; one gains and no one loses, win-neutral; no gain and one loses, neutral-loss; and trades in which there are gains and losses but no way of telling what the net impact is. Since people are rational and self-interested, they would only make the first two types of trade ('Pareto superior trades') and when these are exhausted, there is an equilibrium in which any further trade would lower the general welfare so no further trades occur, an equilibrium state which became known as *Pareto optimal*. This seemed to be a mathematical vindication of Smith's belief that the invisible hand produced the best social outcome and Pareto believed that, 'The theory of economic science thus acquires the rigour of rational mathematics' (though Pareto did not regard his work as vindicating free markets). The work of these three was synthesised by various people including Alfred Marshall and John Hicks' influential textbook *Value and Capital.*

(ii) Von Neumann, 'equilibrium', and the invention of Game Theory

'In his Theory of Games and Economic Behavior, von Neumann laid the foundations for a unified view of information theory, economics, evolution, and intelligence, whose implications continue to emerge.' Dyson.

There are some snippets from pre-20[th] Century figures explaining concepts in terms recognisable through the prism of Game Theory. For example, Ampère wrote '*Considerations sur la théorie mathématique du jeu*' in 1802 (and credited Buffon's 1777 essay on 'moral arithmetic'[370]) and Cournot discussed what would later be described as a specific example of a 'Nash equilibrium' *viz* duopoly in 1838.

However, Game Theory really was born with John von Neumann.[371] In December 1926, he presented the paper '*Zur Theorie der Gesellschaftsspiele*' (*On the Theory of Parlour Games,* published in 1928, translated version [here]) while working on the Hilbert Programme (cf. Endnote on Computing) and quantum mechanics. The connection between the Hilbert Programme and the intellectual origins of Game Theory can perhaps first be traced in a 1912 lecture by one of the world's leading mathematicians and founders of modern set theory, Zermelo, titled '*On the Application of Set Theory to Chess*' which stated of its purpose:

> '… it is not dealing with the practical method for games, but rather is simply giving an answer to the following question: can the value of a particular feasible position in a game for one of the players be mathematically and objectively decided, or can it at least be defined without resorting to more subjective psychological concepts?'

He presented a theorem that chess is strictly determined: that is, either (i) white can force a win, or (ii) black can force a win, or (iii) both sides can force at least a draw. Which of these is the *actual solution* to chess remains unknown.[372]

Von Neumman later wrote:

> '[I]f the theory of Chess were really fully known there would be nothing left to play. The theory would show which of the three possibilities … actually holds, and accordingly the play would be decided before it starts… But our proof, which guarantees the validity of one (and only one) of these three alternatives, gives no practically usable method to determine the true one. This relative, human difficulty necessitates the use of those incomplete, heuristic methods of playing, which constitute 'good' Chess; and without it there would be no element of 'struggle' and 'surprise' in that game.' (p.125)[373]

---

[370] Buffon figured out most of the elements that Darwin would later harmonise in his theory of evolution (Dyson).

[371] The French mathematician Emile Borel also made contributions to early ideas.

[372] Cf. '*Zermelo and the Early History of Game Theory*', by Schwalbe & Walker (1997), which argues that modern literature is full of errors about this paper. According to Leonard (2006), Zermelo's paper was part of a general interest in the game of chess among intellectuals in the first third of the 20[th] century. Lasker (world chess champion 1897–1921) knew Zermelo and both were taught by Hilbert. During the 1920's, there were great debates concerning logical and psychological theories of chess that engaged the intelligentsia beyond mathematicians (eg. Nabokov's *The Defence*).

[373] Elsewhere, he said: 'Chess is not a game. Chess is a well-defined computation. You may not be able to work out the answers, but in theory there must be a solution, a right procedure in any position. Now, real games are not like that at all. Real life is not like that. Real life consists of bluffing, of little tactics of deception, of asking yourself what is the other man going to think I mean to do. And that is what games are about in my theory.'

Von Neumman's 1928 paper proved that *there is a rational solution to every two-person zero-sum game.* That is, in a rigorously defined game with precise payoffs, there is a mathematically rational strategy for both sides - an outcome which both parties cannot hope to improve upon. This introduced the concept of the *minimax*: choose a strategy that *mini*mises the possible *max*imum loss.

Zero-sum games are those where the payoffs 'sum' to zero. For example, chess or Go are zero-sum games because the gain (+1) and the loss (-1) sum to zero; one person's win is another's loss. The famous Prisoners' Dilemma (see below) is a *non-zero-sum* game because the payoffs do not sum to zero: it is possible for both players to make gains. In some games the payoffs to the players are *symmetrical* (e.g. Prisoners' Dilemma); in others, the payoffs are *asymmetrical* (e.g. the Dictator or Ultimatum games). Sometimes the strategies could be completely stated without the need for probabilities ('pure' strategies); sometimes, probabilities have to be assigned for particular actions ('mixed' strategies).

While the optimal minimax strategy might be a 'pure' strategy, von Neumann showed it would often have to be a 'mixed strategy' and this means a return of probability, even if the game itself does not involve probability.

> 'Although … chance was eliminated from the games of strategy under consideration (by introducing expected values and eliminating 'draws'), it has now made a spontaneous reappearance. Even if the rules of the game do not contain any elements of 'hazard' … in specifying the rules of behaviour for the players it becomes imperative to reconsider the element of 'hazard'. The dependence on chance (the 'statistical' element) is such an intrinsic part of the game itself (if not of the world) that there is no need to introduce it artificially by way of the rules of the game itself: even if the formal rules contain no trace of it, it still will assert itself.'

In 1932, he gave a lecture titled '*On Certain Equations of Economics and A Generalization of Brouwer's Fixed-Point Theorem*'. It was published in German in 1938 but not in English until 1945 when it was published as '*A Model of General Economic Equilibrium*'. This paper developed what is sometimes called von Neumann's *Expanding Economic Model* and has been described as the most influential article in mathematical economics. It introduced the use of 'fixed-point theorems':[374]

> 'The mathematical proof is possible only by means of a generalisation of Brouwer's Fix-Point Theorem, i.e. by the use of very fundamental *topological* facts… The connection with topology may be very surprising at first, but the author thinks that it is natural in problems of this kind. The immediate reason for this is the occurrence of a certain 'minimum-maximum' problem… It is closely related to another problem occurring in the theory of games.'

Von Neumann's application of this topological proof to economics was very influential in post-war mathematical economics and in particular was used by Arrow and Debreu in their seminal 1954 paper on general equilibrium, perhaps *the* central paper in modern traditional economics.

In the late 1930's, von Neumann, based at the IAS with Gödel and Einstein, met up with the economist Oskar Morgenstern who was deeply dissatisfied with the state of economics. In 1940, von Neumann began his collaboration on games with Morgenstern, while working on war business including the Manhattan Project and computers, that became *The Theory of Games and Economic*

---

[374] Brouwer's 'fixed point theorem' in topology proved that, in crude terms, if you lay a map of the US on the ground anywhere in the US, one point on the map will lie precisely over the point it represents on the ground beneath. (Buchanan, 2013)

Behavior (TGEB). By December 1942, he had finished his work on this (it was not published until 1944).[375]

TGEB extended the 1928 results and provided a rigorous mathematical method to model economic transactions with a formal theory of games involving strategic interaction. It described the errors that economists had been making in their attitude to the application of mathematics and suggested a different vision. It introduced new mathematical techniques into economics. It provided a different way to examine the concept of equilibrium including 'saddle points'.[376]

TGEB extended the treatment of two-person-zero-sum games to *n-person* games in which coalitions featured among the multiple players.

> 'We shall primarily construct a theory of the zero-sum games, but it will be found possible to dispose, with its help, of all games, without restriction. Precisely: We shall show that the general (hence in particular the variable sum) *n*-person game can be reduced to a zero-sum *n* +1-person game. Now the theory of the zero-sum *n*-person game will be based on the special case of the zero-sum two-person game.'

For his views on the problems of economics at the time and the proper role of mathematics in economics, see the passages quoted at the end of Section 1. After his work on games, von Neumann paid relatively little attention to economics in the next decade, spending his time on maths, physics, weapons, and computers; weather, more than economics, interested him as a nonlinear system ripe for forecasting (cf. Endnote on Computing). In December 1955, not long before he died, he gave a talk on economics which was partially published as '*The Impact of Recent Developments in Science on the Economy and Economics*'.

> 'It is frequently said that economics is not penetrable by rigorous scientific analysis because one cannot experiment freely. One should remember that the natural sciences originated with astronomy, where the mathematical method was first applied with overwhelming success. Of all the sciences, astronomy is the one in which you can least experiment. So the ability to experiment freely is not an absolute necessity. Experimentation is a convenient tool, but large bodies have been developed without it.

> 'It is frequently said that in economics one can never get a statistical sample large enough to build on. Instead, time series are interrupted, altered by gradual or abrupt changes of conditions etc. However, … one realizes that in scientific research as well, there is always some heterogeneity in the material and that one can never be quite sure whether this heterogeneity is essential. The decisive insights in astronomy were actually derived from a very small sample: The known planets, the sun, and the moon.

> 'What seems to be exceedingly difficult in economics is the definition of categories. If you want to know the effects of the production of coal on the general price level, the difficulty is not so much to determine the price level or to determine how much coal has been produced, but to tell how you want to define the level and whether you mean coal, all fuels,

---

[375] Cf. Leonard (1995). Morgenstern later said that von Neumann had urged for their names to be listed alphabetically but that he had refused. He described, probably overly-modestly, his own 'minimal share' but rightly described himself as 'a kind of catalytic factor'.

[376] When the two players' maximin and minimax payoffs match, there is a *saddle point*. Like travellers going through a pass, there is a *maximum* elevation for the person travelling through the pass, and a *minimum* elevation for the person travelling the crest of the pass, amounting to a sort of equilibrium. A saddle point does not guarantee absolute happiness with the outcome, only relative happiness. Many games have payoffs such that there is no saddle point.

or something in between. In other words, it is always in the conceptual area that the lack of exactness lies. Now all science started like this, and economics, as a science, is only a few hundred years old. The natural sciences were more than a millennia old when the first really important progress was made.

'The chances are that the methods in economic science are quite good, and no worse than they were in other fields. But we will still require a great deal of research to develop the essential concepts – the really usable ideas. I think it is in the lack of quite sharply defined concepts that the main difficulty lies, and not in any intrinsic difference between the fields of economics and other sciences.'

Although most economists initially ignored von Neumann's ideas on game theory,[377] a small but influential number, including mathematicians at the RAND Corporation (the first recognisable modern 'think tank'?) led by John Williams, applied it to nuclear strategy as well as economics.[378]

In the 1990's, the movie *A Beautiful Mind* brought John Nash into pop culture, giving the misleading impression that he was the principle developer of Game Theory. Nash's fame rests principally on work he did in 1950-1 that became known as 'the Nash Equilibrium'. In *Non-Cooperative Games* (1950), he wrote:

'[TGEB] contains a theory of *n*-person games of a type which we would call cooperative. This theory is based on an analysis of the interrelationships of the various coalitions which can be formed by the players of the game. Our theory, in contradistinction, is based on the *absence* of coalitions in that it is assumed each participant acts independently, without collaboration or communication with any of the others… [I have proved] that a finite non-cooperative game always has at least one equilibrium point.'[379]

In 1949-50, Merrill Flood, another RAND researcher, began experimenting with staff at RAND (and his own children) playing various games. Nash's results prompted Flood to create what became known as the 'Prisoners' Dilemma' game, the most famous and studied game in Game Theory. It was initially known as 'a non-cooperative pair' and the name 'Prisoners' Dilemma' was given it by Tucker later in 1950 when he had to think of a way of explaining the concept to his psychology class at Stanford and hit on an anecdote putting the payoff matrix in the form of two prisoners in separate cells considering the pros and cons of ratting on each other. Prisoners' Dilemma has been called 'the *E. coli* of social psychology' by Axelrod, so popular has it become in so many different fields. Many studies of *Iterated Prisoners' Dilemma* games have shown that generally neither human

---

[377] Martin Shubik, a Princeton mathematician, recounted the scene he found: 'The contrast of attitudes between the economics department and mathematics department was stamped on my mind… The former projected an atmosphere of dull-business-as-usual conservatism… The latter was electric with ideas… When von Neumann gave his seminar on his growth model, with a few exceptions, the serried ranks of Princeton economists could scarce forebear to yawn.'

[378] E.g Albert Wohlstetter published his *Selection and Use of Strategic Air Bases* (RAND, R-266, sometimes referred to as *The Basing Study*) in 1954. Williams persuaded the RAND Board and the infamous SAC General Curtis LeMay to develop a social science division at RAND that could include economists and psychologists to explore the practical potential of Game Theory further. He also hired von Neumann as a consultant; when the latter said he was too busy, Williams told him he only wanted the time it took von Neumann to shave in the morning. Kubrick's *Dr Strangelove* satirised RAND's use of game theory.

[379] Von Neumann apparently remarked of Nash's results, 'That's trivial you know. It's just a fixed point theorem.' Nash himself said that von Neumann was a 'European gentleman' but was not keen on his results.

nor evolved genetic algorithm players converge on the Nash equilibrium but choose to cooperate far more than Nash's theory predicts (cf. Section 7 and footnotes re recent breakthroughs).[380]

Von Neumann's brief forays into economics were very much a minor sideline for him but there is no doubt of his influence. Despite von Neumann's reservations about neoclassical economics, Paul Samuelson admitted that, 'He darted briefly into our domain, and it has never been the same since.'

<div align="center">*</div>

(iii)   General Equilibrium Theory and the Neoclassical Synthesis...

The 'Neoclassical synthesis' of the marginal revolution and what became known as General Equilibrium Theory (GET) is usually defined as resting upon two economists in particular, Paul Samuelson and Kenneth Arrow.

Although utility had become a central feature of economics, it was clearly not an objective value like energy. Samuelson argued, in *Foundations of Economic Analysis* (1941),[381] that relative utility was revealed by rational optimisers displaying their preferences through choices exercised in a market. Samuelson set out a formal mathematical approach to economic analysis and *equilibrium*. Decades later, he described his explicit application of formal mathematics and the techniques of thermodynamics:

> 'Perhaps most relevant of all for the genesis of *Foundations*, Edwin Bidwell Wilson (1879–1964) was at Harvard. Wilson was the great Willard Gibbs's last (and, essentially only) protégée at Yale. He was a mathematician, a mathematical physicist, a mathematical statistician, a mathematical economist, a polymath who had done first-class work in many fields of the natural and social sciences. I was perhaps his only disciple... I was vaccinated early to understand that economics and physics could share the same formal mathematical theorems (Euler's theorem on homogeneous functions, Weierstrass's theorems on constrained maxima, Jacobi determinant identities underlying Le Chatelier reactions, etc.), while still not resting on the same empirical foundations and certainties.'

In a footnote to a discussion of comparative static analysis in *Foundations*, he wrote that 'this is essentially the same method of thermodynamics'.

In 1954, Arrow and Debreu published a theorem which was one of the most influential post-war ideas in economics. They used Brouwer's fixed-point theorem (as von Neumann had done twenty years before) to prove what Walras had wanted to prove - a set of prices must exist that coordinate everyone's choices and produce an equilibrium.

> '[They] connected Walras' notion of general equilibrium with Pareto's concept of optimality in a very general way, thus creating the Neoclassical theory of general equilibrium. Their theorem showed that all the markets in the economy together would automatically coordinate on a set of prices that was Pareto optimal for the economy as a whole, and that this would occur even when there was uncertainty in the market (Walras had required …

---

[380] The game was discussed and played at RAND without publishing. Flood wrote up the results in 1952 as an internal RAND memo accompanied by the real-time comments of the players. In 1958, Flood published the results formally (*Some Experimental Games*). Flood concluded that 'there was no tendency to seek as the final solution … the Nash equilibrium point.' Thousands of subsequent experiments found the same.

[381] This was Samuelson's PhD thesis and, as Lo says, would have been a presumptuous title but for the fact that it did actually become the foundation of economic analysis. It was published as a book in 1947.

that everything be certain). This automatic coordination occurs because markets are linked with each other by the ability of some goods to act as *substitutes* for others … and by the tendency of other goods to be consumed together as *complements*… Arrow and Debreu showed that prices act like a nervous system, transmitting signals about supply and demand throughout the economy, and that self-interested people react to those price signals and inevitably drive the system to its socially optimal equilibrium point.' (Beinhocker, p. 38).

The basic features of the general equilibrium theory were:

- Agents are assumed to have rational preferences.
- Agents are assumed to have complete information about the probabilities of future events.
- Agents are assumed to 'maximise utility'.
- The standard model for pricing states, 'Set the price equal to the marginal cost', where the 'marginal cost' means the cost of producing an additional unit of output.
- Prices will bring supply and demand into equilibrium.
- Futures markets exist for every product and service.[382]

Debreu later wrote:

'Before the contemporary period of the past five decades, theoretical physics had been an inaccessible ideal toward which economic theory sometimes strove. During that period, this striving became a powerful stimulus in the mathematization of economic theory.'

The '*neoclassical synthesis*' refers to the synthesis of the ideas of Walras, Jevons, Samuelson, Arrow *et al* with the ideas of Keynes. It was given particular force by Samuelson's textbook, *Economics: An Introductory Analysis* (first edition 1948), possibly the most influential economics textbook written.

However, while Arrow-Debrau had proved that, given their assumptions, an equilibrium must exist, they had *not* proved that their equilibrium must be *stable*. In the 1970's, Debreu, Mantel, and Sonnenschein showed that a plausible economy may never reach a stable Arrow-Debreu equilibrium and since then *nobody has proved that even very simplified versions of the Arrow-Debreu model (never mind a plausible model of the real economy) must have a stable equilibrium*. To prove the existence of a stable equilibrium, economists have been forced to make ludicrous assumptions, and, Buchanan argues, this is why research in the area largely ended and orthodox economics ignored this embarrassing issue and embedded the Arrow-Debreu model regardless. However, the lack of a stable equilibrium should be no surprise; as Donald Saari notes, Walras' equations are 'sufficiently elementary to be taught in a first course on vector calculus. So, we must wonder what a nice, simple model is doing in a complex place like this.'

One of the reasons why economists were comfortable with this approach is an extremely influential paper by Milton Friedman. Friedman was arguably right about a lot of things concerning the superiority of markets over central planning but his 1953 paper '*The Methodology of Positive Economics*' was a disaster. He correctly argued that economics should, like physics, test its hypotheses by comparison with data. He then argued that a hypothesis should *not* be judged by the realism of its assumptions: the inability of the economist to do experiments like a natural scientist 'makes it tempting to suppose that … the conformity of these "assumptions" [the hypothesis'] to

---

[382] Although these mathematical models purported to demonstrate how markets could work optimally, many of the economists involved in these debates from the late 19th Century argued, like Pareto, that central planners could solve the equations just as well as markets. The Marginalist models were not, therefore, necessarily great arguments for market economies. E.g. because Schumpeter accepted general equilibrium theory, he argued that 'there is nothing wrong with the pure logic of socialism'.

"reality" is a test of the validity of the hypothesis… This widely held view is fundamentally wrong and productive of much mischief.' There is some sense in this: all models simplify, by definition, so one cannot dismiss a model simply because it is based on assumptions that are unrealistic. However, Friedman then wrote, 'Truly important and significant hypotheses will be found to have "assumptions" that are wildly inaccurate descriptive representations of reality, and, in general, the more significant the theory, the more unrealistic the assumptions… To be important, therefore, a hypothesis must be descriptively false in its assumptions.' This argument has been used to justify many an economic theory, including Sharpe's extremely influential Capital Asset Pricing model (see below). Physics would be as useful as economics if it operated on the basis that 'the more significant the theory, the more unrealistic the assumptions'.

General equilibrium theory stresses stability, order, and equilibrium. It remained the dominant paradigm for decades *not* because it obviously embodied a realistic psychological theory about the world, *nor* because the epistemological assumptions of perfect information and static equilibrium make sense, which they do not, but because the model allowed deduction from axioms and provided a veneer of physics's rigour.

*

(iv)   The birth of Orthodox Financial Theory (OFT) in a 'Random Walk'…

Bachelier and the Random Walk…

In the 19th Century, the French mathematician Fourier had devised equations describing the way heat spreads. Another French mathematician, Louis Bachelier, applied Fourier's equations to the probability of prices moving up or down. In 1900, he wrote a dissertation called '*The Theory of Speculation*' in which he claimed that financial market prices move in a *random walk* according to a normal distribution like the random movement of molecules in a gas or liquid. Decades earlier, Robert Brown had observed the random movement of pollen grains – '*Brownian motion*'. In 1905, five years after Bachelier's thesis, Einstein famously developed equations for Brownian motion which were 'very similar' (Mandelbrot) to Bachelier's equations for bond price probability, though Einstein did not know this since Bachelier's thesis would not become known until after 1945. 'One cannot help but marvel that the movement of security prices, the motion of molecules, and the diffusion of heat could all be of the same mathematical species' (Mandelbrot).

What does the *random walk* mean? If you see a drunk staggering across an open field where will he go? A few steps here, a few steps there but on average he gets nowhere.

> 'So if you consider only that average, his random walk across the field will be forever stuck at his starting point. And that would be the best possible forecast of his future position at any time, if you had to make such a guess.' (Mandelbrot)

The same logic applies to prices. In the absence of new information, the best forecast of tomorrow's price is *today's price*. Further, each variation is unrelated to the last.

> 'The price changes, in the language of statistics, form a sequence of independent and identically distributed random variables.' (Mandelbrot)

Bachelier argued, therefore, that one could plot prices over time, see a bell curve, and apply the mathematical tools used for normal distribution, though his theory also implied that the mathematical expectation of the speculator is zero.

However, Bachelier's ideas did not catch on and he was ignored for decades until his ideas were revived after 1945 and used as the foundation of modern financial theory.

> '[Bachelier's] work marks the twin births of the continuous-time mathematics of stochastic processes and the continuous-time economics of option pricing. In analyzing the problem of option pricing, Bachelier provides two different derivations of the Fourier partial differential equation as the equation for the probability density of what is now known as a Wiener process/Brownian motion. In one of the derivations, he writes down what is now commonly called the Chapman-Kolmogorov convolution probability integral, which is surely among the earlier appearances of that integral in print. In the other derivation, he takes the limit of a discrete-time binomial process to derive the continuous-time transition probabilities. Along the way, Bachelier also developed essentially the method of images (reflection) to solve for the probability function of a diffusion process with an absorbing barrier. This all took place five years before Einstein's discovery of these same equations in his famous mathematical theory of Brownian motion.' (Lo & Merton, 2009).

Post-war: 'the Efficient Market Thesis' is born…

Although the 1929 Crash and Great Depression triggered a new interest in the theory of markets, Bachelier's thesis remained in obscurity. In 1955, Leonard Savage (who had worked with von Neumann during the war) discovered it and sent a note about it to various people including Paul Samuelson. In 1965, Paul Samuelson published '*Proof that Properly Anticipated Prices Fluctuate Randomly*' which showed that a market with rational investors using all available information *must* be unpredictable - a predictable market is self-contradictory. Samuelson's paper was mathematical - he stressed he was not talking about real markets. But the idea provoked a few people to consider its implications for real markets.

Most influentially, Eugene Fama published *Efficient Capital Markets* in 1970, arguing that *the inability to forecast a market - its apparent randomness - is a flipside of the market's efficiency in processing available information:* 'A market in which prices always "fully reflect" available information is called "efficient."' Fama's thesis became known as '*the Efficient Market Thesis*' (EMT).[383] In 1973, Burton Malkiel, a Princeton professor, wrote a book called *A Random Walk Down Wall Street* which popularised the basic ideas.

The Efficient Market Hypothesis states that the market correctly prices all information; new information is an 'exogenous' factor that arrives like the random toss of the coin, therefore prices move in a random walk with a normal distribution and therefore, by definition, since historical price information is uncorrelated with and independent of future prices, neither historical nor current prices contain information that can be profited from.

> 'In an ideal market, security prices fully reflect all relevant information. A financial market is a fair game in which buyer balances seller. Given that, the price at any particular moment must be the 'right' one. Buyer and seller may differ in opinion; one may be a bear, and another a bull. But they both agree on the price, or there would be no deal. Multiply this thinking by the millions of daily deals of a bustling market, and you conclude that the general market price must be 'right' as well – that is, that the published price reflects the market's overall best guess, given the information to hand, of what a stock is likely to profit its owner. And if that is true … then you cannot beat the market… That does not mean you cannot win, ever. In fact, by the simple odds of a fair game, you can expect to win half the time and lose half the time.

---

[383] Mandelbrot was Fama's thesis adviser but the latter clearly did not agree with the former's ideas on markets.

And if you have special insights into a stock, you could profit from being the first in the market to act on it. But you cannot be sure you are right or first… So, in sum, it may not be worth your while to spend all that time and money getting the information in the first place. Cheaper and safer to ride with the market. Buy a stock index fund. Relax. Be passive.' Mandelbrot

Samuelson concluded:

'A respect for evidence compels me to incline toward the hypothesis that most portfolio decision makers should go out of business – take up plumbing, teach Greek, or help produce the annual GNP by serving as corporate executives.'

However, as Mandelbrot says, Wall Street is nothing if not flexible:

'[W]hat could have been its epitaph was recast as a rallying cry. Bachelier's thesis was elaborated into a mature theory… It came in the 1970s and 1980s to be the guiding principle for many of the standard tools of modern finance, the orthodox line, taught in business schools and shrink-wrapped into financial software packages, for how to value securities, how to build portfolios, how to measure financial performance, and how to judge the merits of a financial project. As will be seen, it is a house built on sand.'

On top of Bachelier's ideas and the emerging EMH, three tools were developed which have remained fundamental for half a century. The standard theory is to value an asset by expected future cash flows combined with a discount rate that reflects risk ('net present value'). However, pricing risk is at the heart of the problem.

1. *Modern Portfolio Theory* (MPT), developed by Harry Markowitz. Markowitz pondered how to combine equations for reward and risk using the maths of the normal distribution. He concluded that a) you work out a forecast price using various tools and b) you use the normal distribution to tell you the historical *variance* which is a measure of volatility; that is, everything could be expressed by the two numbers of 'the mean and variance of what you expect the stock will pay back by the time you sell'. With your forecast price and the bell curve for past variance, you have a method to predict how much you expect to make with *x* variance. Now you can compare many assets and build a portfolio that supposedly combines different kinds of assets that will respond to events in different ways; that is, they should be *uncorrelated* like tosses of the coin. 'Diversification is both observed and sensible,' declared Markovitz.

'Thus, with Markowitz's math, for each level of risk you contemplate, you can devise an efficient portfolio that will yield the highest possible profit. And for each level of profit you target, there is an efficient portfolio with the lowest possible risk.' (Mandelbrot)

From this concept came the appropriate term 'financial engineering'. Markowitz later recalled the lines from Merchant of Venice:

'I thank my fortune for it,
My ventures are not in one bottom trusted,
Nor to one place; nor is my whole estate
Upon the fortune of this present year;
Therefore my merchandise makes me not sad.'

2. *The Capital Asset Pricing Model*, developed by William Sharpe in the 1960's. Sharpe studied with Markowitz. Sharpe asked the question: what if everybody plays by Markowitz's rules? The answer: everybody would converge on just one portfolio, 'the market portfolio', 'thus was born the notion of a stock-index fund' (Mandelbrot). He also developed a model to price the volatility of a stock relative to the market (known as the stock's '*beta*'.) It enormously simplified Markovitz's system and became known as the *Capital Asset Pricing Model* (CAPM). After his initial paper was published in 1964, there was almost no reaction. Now, billions of dollars depend upon this idea.

3. The *Black-Scholes formula* for pricing options, developed by Fischer Black and Myron Scholes. How can one price an option when one does not know the future price of the underlying asset? Black and Scholes created a formula based upon a) the terms of the option ('strike price' and time to expiry) and b) the volatility. As the option matures and the stock price moves, the option's value will change. The Black-Scholes formula provided a mathematical tool to analyse this change relative to the (supposed) variance. It was also based on the 'random walk' of normal distribution; the 'stock price follows a random walk in continuous time' (Black-Scholes). They also had trouble getting the initial paper published and it was only published in 1973 one month after the opening of the Chicago Board Options Exchange in April.[384]

These three intellectual constructs are fundamental to MBAs and economics degrees and hence to the education of leading businessmen and financiers and they all rely on Bachelier's theoretical foundation.

> 'The whole edifice hung together – provided you assume Bachelier and his latter-day disciples are correct. Variance and standard deviation are good proxies for risk, as Markowitz posited – provided the bell curve correctly describes how prices move. Sharpe's beta and cost-of-capital estimates make sense – provided Markowitz is right and, in turn, Bachelier is right. And Black-Scholes is right – again, provided you assume the bell curve is relevant and that prices move continuously. Taken together, this intellectual edifice is an extraordinary testament to human ingenuity. But the whole is no stronger than its weakest member.' (Mandelbrot)

Problems with traditional economics and financial theory.

1. People are not 'perfectly rational'. There is not even a clear definition of what this term means.

2. People do not have perfect information.

3. People are not homogenous - they are heterogeneous in their beliefs and emotions.

4. Prices do not move according to the normal distribution. Mandelbrot first drew attention to this in the 1960's but was largely ignored: if the Dow Jones Industrial Average price movements *were* normally distributed 1916-2003, there should have been six days of >4.5% movement but there were 366; there should only be a day of >7% movement once every 300,000 years but there were 48 such days (Mandelbrot) 'Patterns of price movements are not random. However, they're close enough to random so that getting some excess, some edge out of it, is not easy and not so obvious' (Jim Simons, CEO of Renaissance, the world's most successful hedge fund). Nor do prices

---

[384] The legendary investor Ed Thorp independently worked out a similar formula before they published which he used to make a fortune. Shortly after publication, Texas Instruments produced a new handheld calculator that could price options using the BSF. Supposedly when Black was introduced to the head of trading at Goldman Sachs, the latter said, 'So you're Fischer Black. Let me tell you something. You don't know shit about options.' Black said after he had joined Goldman Sachs that 'markets look a lot more efficient from the banks of the Charles than from the banks of the Hudson'. Scholes and Merton lost a lot of money and kudos when the hedge fund LTCM blew up in 1998 (*The Quants*, 2010). The Black-Scholes formula for option pricing 'is also the solution to the heat equation' (Lo, 2010).

move continuously: 'Discontinuity, far from being an anomaly best ignored, is an essential ingredient of markets that helps set finance apart from the natural sciences' (Mandelbrot).

There is *still* no agreed explanation for the 22% crash in the Dow in October 1987. There are different ideas - small events cascaded because of 'portfolio insurance' (as some started to sell, this triggered others' insurance systems to trigger computerised selling), and/or because of signs of a 1929 collapse in the days before - but nobody can explain the timing. Similarly, there is still no agreed explanation for the Flash Crash of 6 May 2010.

Some big moves are clearly related to big news (e.g. outbreak of war, Eisenhower's heart attack, 9/11) but others occur in the absence of such news (e.g. 3/9/1946, 4/6/1962, October 1987, 6/5/2010 'the flash crash'). Various studies have shown that *most* of the big movements are *not* connected to big news - *they are self-generated by market dynamics* - and in general prices are more volatile than the EMH predicts based on expected future dividends. As HFT takes over, this seems to be happening more often - Nanex reports many mild versions of the flash crash which are inexplicable in terms of markets responding to new information in the real world (i.e signals outside the HFT system itself). (Buchanan, 2013)

Although we cannot predict either earthquakes or market crashes, it is interesting that they both follow an almost identical mathematical rule: both are distributed in an '*inverse power law*', with the probability of an event being inversely proportional to the size of the event (lots of small small events and few bigger ones). High-frequency algorithmic trading may be both increasing liquidity (for a while anyway) *and* making volatility worse: '*together with the tendency of liquidity to evaporate in dangerous times, and it starts to look like what happens on subsecond times may reflect a fundamental transition as markets break free of human control*' (Buchanan).[385]

5. Price movements are not independent. Volatility clusters: a large movement today makes a large movement tomorrow more likely. Not only are large price movements more likely than a normal distribution predicts but the market has a 'memory' of the size of price movement. Following a big crash, the probability of 'aftershock' events falls off in the same way as for earthquakes, in direct proportion to the time since the crash. Tobias Preis analysed price movements on scales from milliseconds to decades and found that, regardless of scale, there is a signal in markets as momentum is about to shift direction: trading volume increases as the 'switching' moment approaches. Financial markets are, like turbulent physical systems, full of *fractal* patterns (self-similar over different scales): if you remove the labels from axes, you cannot distinguish whether one is looking at days, weeks, months, or years.

6. Markets, like food chains, do not operate at 'equilibrium', but in 'dis-equilibrium', nor do they tend towards a 'stable equilibrium'.

7. Most published economic forecasting of things like future GDP and employment is wrong (see Section 6).
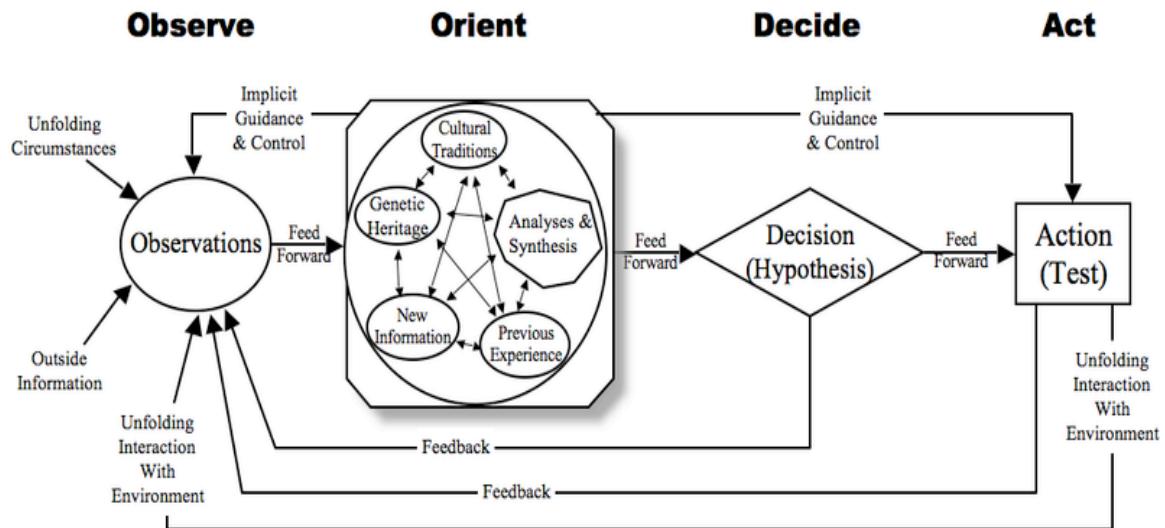
Much of the economics establishment has not changed its tune despite 2008. Robert Lucas said that the EMH had suffered 'a flood of criticism which has served mainly to confirm the accuracy of the hypothesis.' Fama said: '[EMH] did quite well in this episode [2008]… I don't even know what a

---

[385] Doyne Farmer *et al* did an experiment in which games are generated randomly and learning algorithms, starting with random strategies, repeatedly play them modifying their strategies. When the number of strategies is small (<5), the algorithms rapidly converge on a Nash equilibrium but when the number is larger (e.g just 50), they do not: 'Their strategies continually vary as each player responds to past conditions and attempts to do better than the other players. The trajectories in the strategy space display high-dimensional chaos, suggesting that for most intents and purposes the behavior is essentially random, and the future evolution is inherently unpredictable.' (Farmer & Galla).
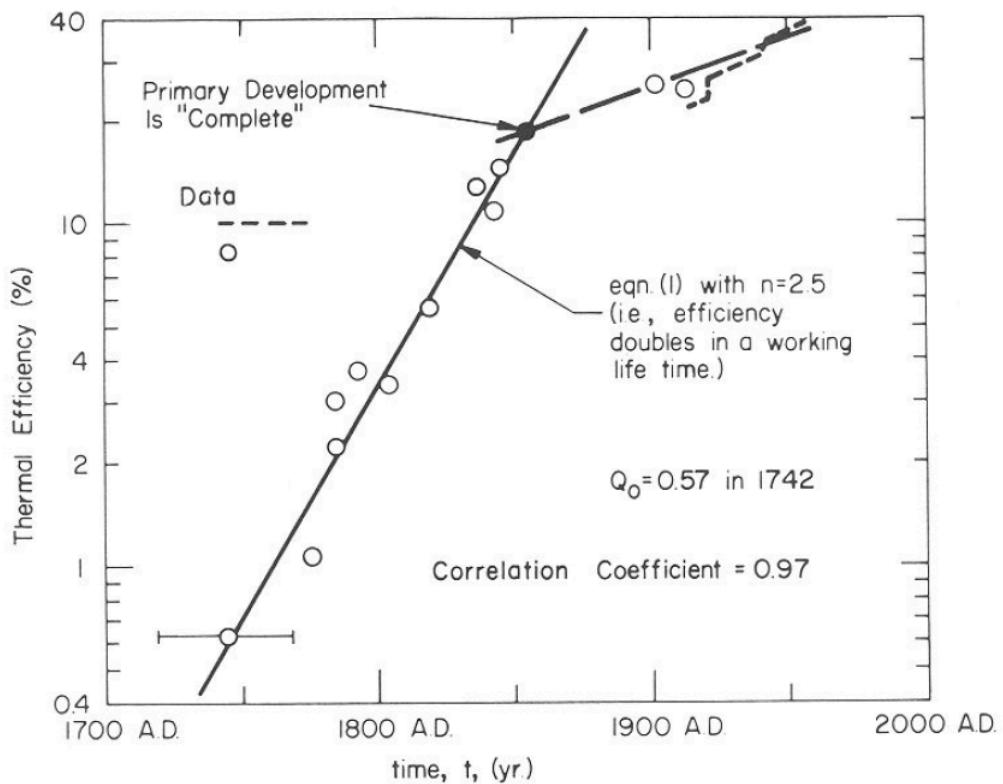
bubble means. These words have become popular. I don't think they have any meaning ... They have to be predictable phenomena. I don't think any of this was particularly predictable.' As Buchanan says, it's as if Fama's line at a geophysicists conference would be 'ignore earthquakes because we can't predict them'.

Endnote: Boyd's 'OODA loop'
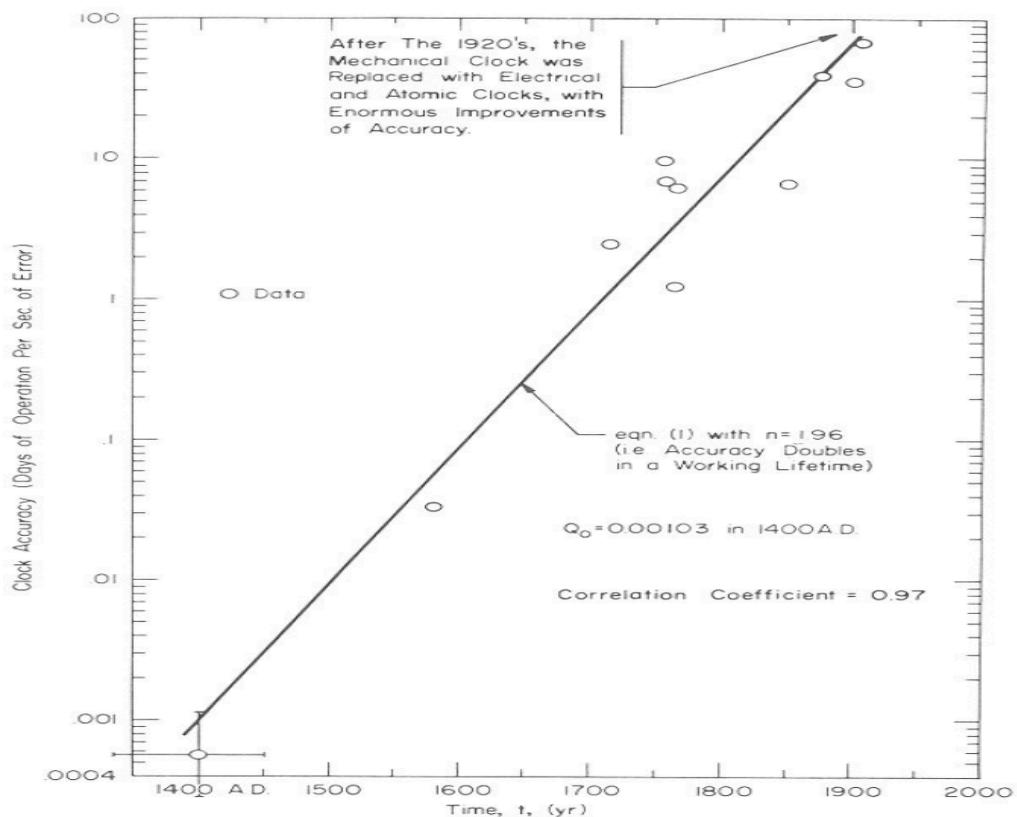
Endnote: Some interesting charts
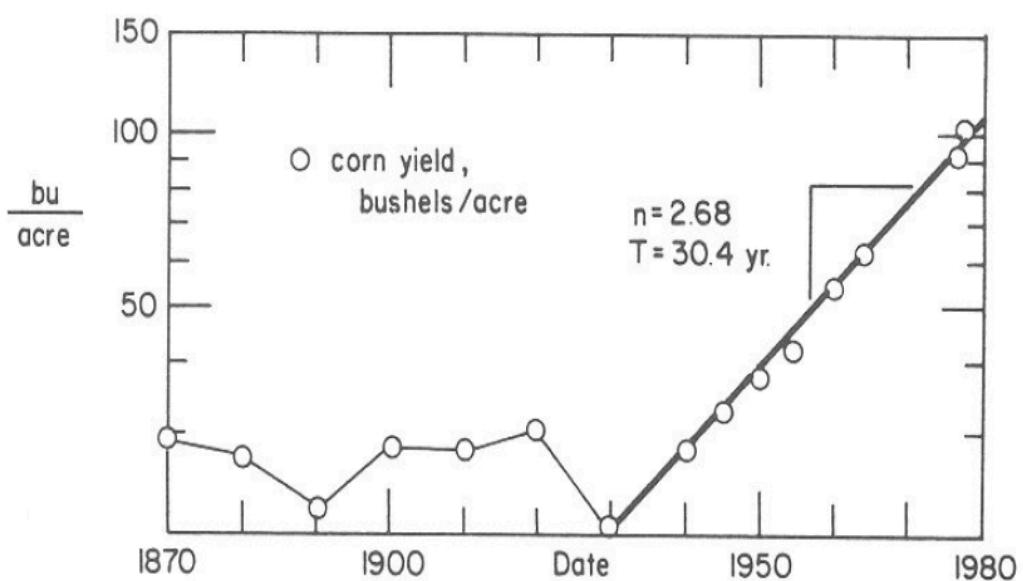
Increase of thermal efficiency with time for hydrocarbon-fueled steam power plants (Lienhard, 1985)
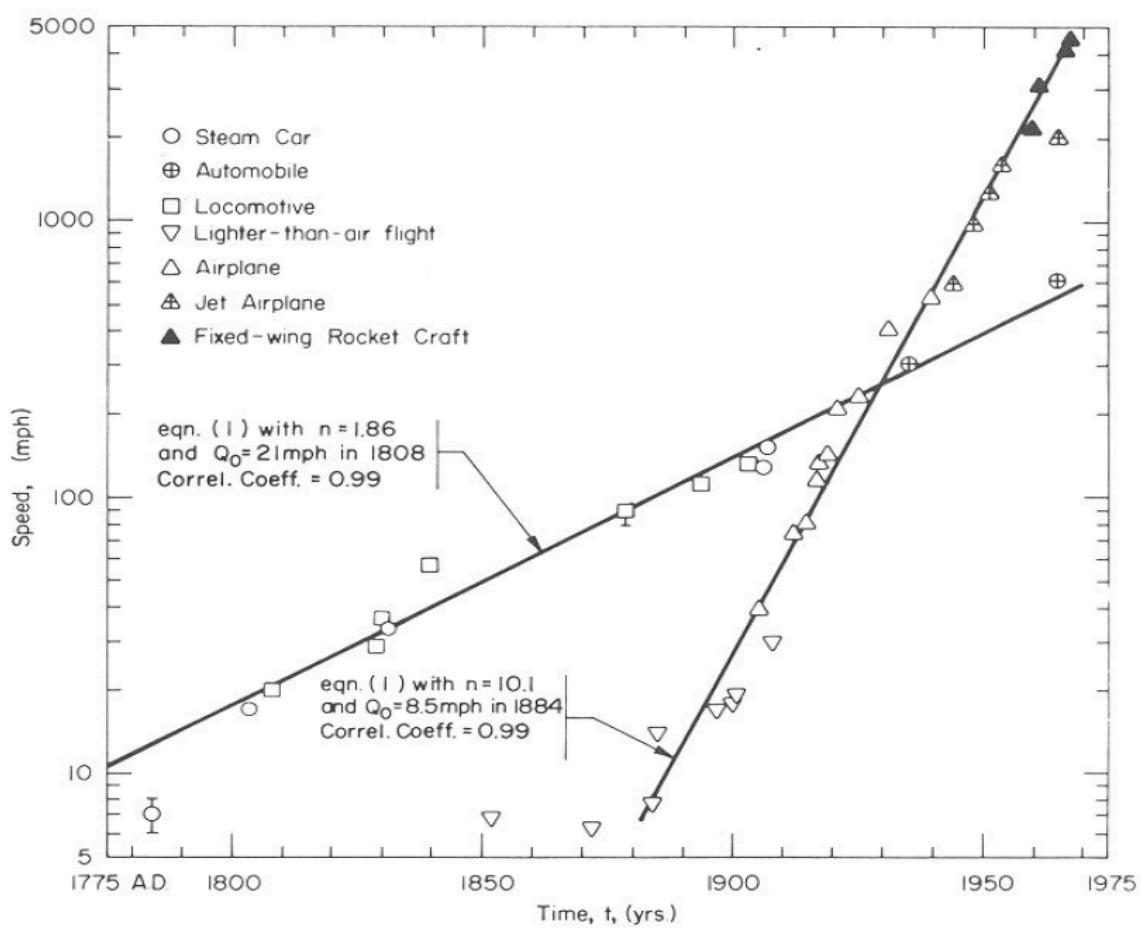
Increase of mechanical clock accuracy with time (Lienhard)
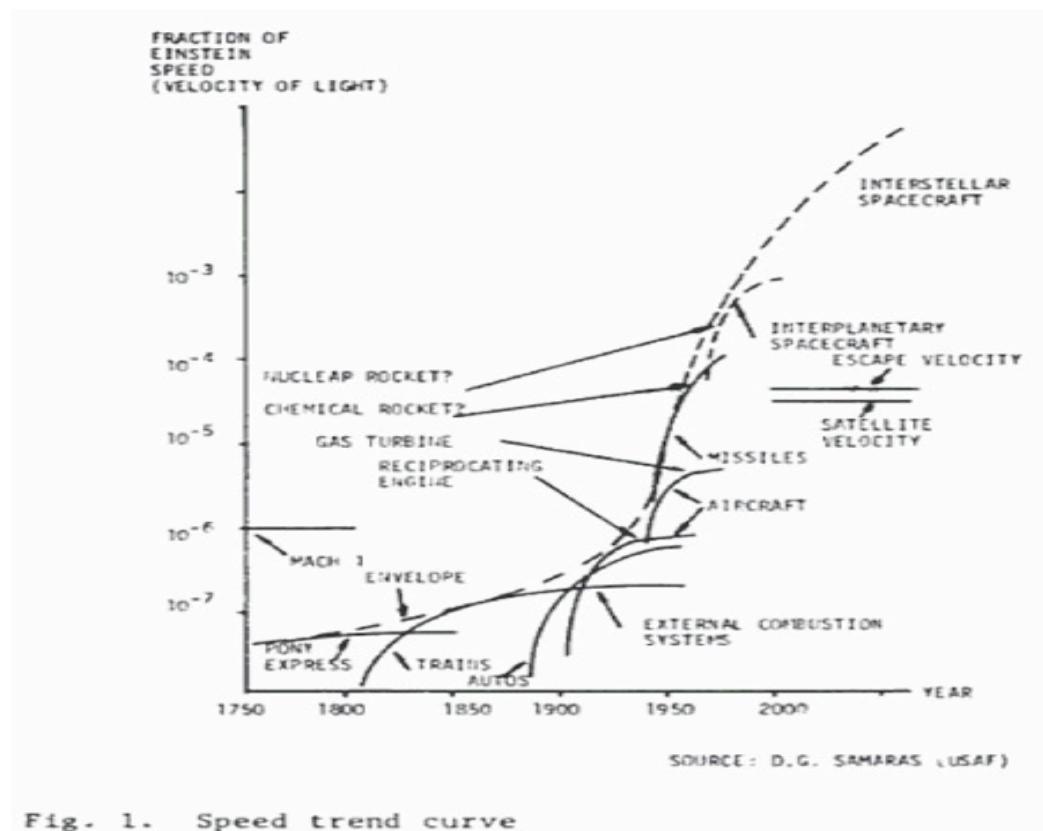
Increase of corn yields with time (Lienhard)

Increase of speed of transportation with time (Lienhard)
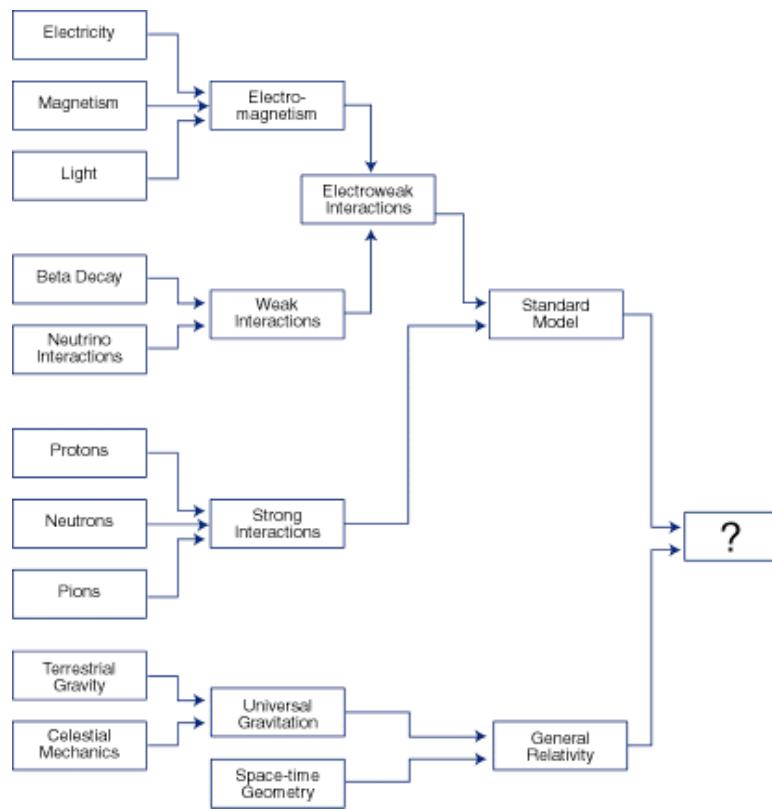
Increase of speed of transportation with time (United States Air Force, 1953)

Fig. 1.   Speed trend curve

Doubling times of various technological performance in months (Kelly)

Fiber optic throughput	wavelengths per fiber	9
Optical network	$/bit	9
Wireless	bits per second	10
Communication	bits per dollar	12
Magnetic areal storage	gigabit/in2	12
digital cameras	pixels per dollar	12
Microprocessor	$ per cycle	13
Supercomputer power	flops	14
RAM	MiB/$	16
RAM	bits per dollar	18
DNA sequencing	$ per base pair	18
Transistor	$ per transistor	18
PCU Power consumption	watts/cm2	18
Pixels	per array	19
Harddrive Storage	Gigabyte per $	20
Chip	MIPS	21
DNA sequencing	$ per base pair	22
Trunk line data speeds	bits/sec	22
Microprocessor	transistors per	24
Chip processor	MHz/$	27
Bandwidth	kilobits per second per $	30
Microprocessor	hrtz	36

Chart:The development of physics (Weinberg)

Endnote: Some ideas for an Odyssean Reading List

Maths and Computer Science

There are some basic areas that are necessary to understand the foundations of pure and applied maths and physics, including (A Level modules in brackets):

- Algebraic geometry and trigonometry (GCSE / C1).
- Sequences and series (C2).
- Logarithms, exponents and exponential functions (C2).
- Combinations and Permutations (S1).
- The Law of Large Numbers, Normal Distributions, and the Central Limit Theorem (S2).
- Conditional Probability (S1) and Bayes' Theory (S2).
- Correlation and Hypothesis testing (S1 & S2).
- Differentiation (C1), Integration (C2), and Differential Equations (C4).
- Complex numbers (FP1).
- Linear Algebra, Matrices (FP1).
- Group Theory (FP3).
- Set theory, logic, and proof. (Not in Maths or Further Maths A Level.)

There are various outstanding books.

Some textbooks recommended by leading professors: *The Method of Coordinates*, Gelfand, Glagoleva, Kirillov. *Functions and Graphs*, Gelfand and Glagoleva. *Lines and Curves*, Vassiliev and Gutenmacher. (These were written for the Russian 'correspondence school' - a way of giving talented maths pupils a useful curriculum in such a vast country. Many who used them went on to the famous Kolmogorov schools.)

Mathematics for the Nonmathematician, Morris Kline. A reader with no more than GCSE Maths can read this introduction to maths from Greece through the birth of calculus.

Mathematics and the Physical World, Morris Kline.

Innumeracy, John Allen Paulos (1988). A classic, recommended by Tim Gowers and others, for a non-specialist introduction to mathematical reasoning.

Reckoning With Risk, Gigerenzer. On typical problems dealing with statistics, Bayes' Theory, and how to improve understanding of probability.

Mathematics: A short introduction, Gowers. For the beginner, by a world-leading mathematician.

The Story of Mathematics, Marcus du Sautoy. Book accompanying an OU course. Very good introduction to some fundamentals, from Pythagoras to Newton to e and complex numbers.

Symmetry, Marcus du Sautoy.

Number, Dantzig (1930, updated 1953; new edition 2007). A non-specialist can enjoy this though much will be opaque.

The Joy of X, Strogatz. This 2012 book, based on a series of NYT blogs, introduces many areas of maths, such as complex numbers.

A good online course is Scott Page's '_Model Thinking_' run by Coursera. It does not require A Level Maths but it is a great introduction to many of the topics covered in this essay (e.g an introduction to some game theory, network theory, information theory, basic economic ideas, conditional probability, correlations, various kinds of computer model etc) and it introduces a lot of quite advanced concepts not normally taught in schools but often referred to in non-school texts that it is very handy to have a feel for (e.g. fitness landscapes, Markov processes).

More advanced…

How to Solve It, Polya. A classic book on mathematical problem solving.

Solving Mathematical Problems, Terence Tao. A modern version of Polya for advanced pupils, by a Fields Medalist.

Causality, Judea Pearl.

What is mathematics?, Courant. A general introduction to number theory, topology, calculus and other subjects. Partly accessible to a non-specialist with some maths already, though far from easy (blurb endorsement by Einstein).

Gödel's Theorem: An Incomplete Guide to its Use and Abuse, Franzen. The best book about Gödel's Theorem (according to the editor of Gödel's Collected Works) which explains why almost everything one reads about it - including by some famous scientists (e.g. Dyson, Hawking) - is wrong.

Mathematics: Its content, methods and meaning, Kolmogorov _et al._

Mathematical Thought from Ancient to Modern Times (_3 volumes_), Morris Kline. This is a scholarly history of maths, not for a general reader.

Calculus, Spivak (2008 edition). A classic book on calculus / analysis designed for undergraduates.

Hardy's A mathematician's apology is cited by many professional mathematicians as an inspiration.

Mathematical Education, Bill Thurston.

Tim Gowers' blog

Terence Tao's blog. E.g. Tao on sets and logic.

On Computer Science… One can imagine many pupils preferring to study the online course of _Artificial Intelligence_ (Russell & Norvig, 3rd edition, 2009) than any existing options available to them 16-18. Both Sebastian Thrun (Udacity) and Andrew Ng (Coursera) now offer online courses introducing machine intelligence. It would be a great step forward for schools to replace many of their traditional courses with pupils _and teachers_ doing these courses, and providing the necessary background maths. Scott Aaronson, one of the leading authorities on quantum computers, has begun an interesting course that teaches some basics of computer science and computational complexity theory. _Wired for War_, Peter Singer (for the general reader, on robots, AI and war). Gleick has written a history of information theory and computation, and Dyson the story of Turing, von Neumann and the computer.

Introduction to the Theory of Computation, Sipser (2005 edition). Classic text, university level, partly accessible to someone doing A Level Maths.

On *Physics…*

Six Easy Pieces, Feynman. A classic based on his famous lectures. A short companion book to his famous lecture series, *Tips on Physics: a problem-solving supplement*, which was unpublished for many years, is useful for someone doing A Levels or beginning an undergraduate course. Here is a wonderful interview with Feynman.

Physics for Future Presidents, Richard Muller (2010). Brilliant physics book for the interested non-specialist written by a top physicist, widely praised by Nobel Prize winners, used in his Berkeley course voted 'best course on campus'.

Theoretical Minimum, Leonard Susskind (2013). Based on a course Susskind taught in San Francisco to give people a basic physics education (maybe wait until the typos are fixed).

Fab: From Personal Computers to Personal Fabrication, Gershenfeld (MIT). On digital fabrication.

The Great Equations, Crease.

A Cavendish Quantum Mechanics Primer, Professor Mark Warner. This excellent short book introduces quantum mechanics using A Level maths.

Histories of the Standard Model: *The Second Creation*, Crease & Mann (1996) and *The Infinity Puzzle*, Frank Close (2011). A great biography of Einstein by Isaacson. A great biography of Dirac, *The Strangest Man*, Farmelo.

The Dance of the Photons, Zeilinger. On the connections between quantum physics, computation and information theory.

The Limits of Quantum Computers, Scott Aaronson (Scientific American, 2008). *Quantum Computing Since Democritus*, Scott Aaronson (2013) is a brilliant introduction to many ideas about computation, physics and the frontiers of quantum information science. Michael Nielsen on quantum mechanics and computers: *Why the world needs quantum mechanics… Quantum Computing for Everyone… Quantum Computing for the Determined*.

Woit's blog… Dorigo's blog. Great for keeping up with developments in physics.

On genetics, the Brain, Mind, AI…

This review (*Natural selection 150 years on*, Pagel, Nature, 2009) summarises recent developments in evolutionary thinking. On evolutionary biology and psychology, behavioural genetics, cognitive science, neuroscience… Pinker's *The Blank Slate*. Plomin's *Behavioral Genetics* is a great standard textbook. *What is Thought?*, Baum. Consciousness, Koch. *Thinking, fast and slow*, Kahneman (cf. Kahneman on cognitive biases and foreign policy). *Expert Political Judgment*, Tetlock. *Sources of Power*, Klein (on expert decisions under pressure). Drexler (nanotech pioneer) describes how he educates himself. Hamming on how to do research. On *biological engineering* and *synthetic biology*: Carlson's *Biology is Technology*. On the future of the genetics revolution, *The Language of Life* by Francis Collins, former director of the Human Genome Project. (Genetics has transformed the

picture of pre-history by allowing analysis of migration patterns. It is also challenging historical analysis of modern events such as the Industrial Revolution. Although it has been thought until recently that we are largely the same psychologically as the ancestral population that migrated from Africa ~50,000 years ago, many recent studies have shown that mutations have changed features (e.g. altitude tolerance, lactose tolerance) in much shorter periods than previously thought. Historians have looked for evidence of such changes. E.g. Clark argues, controversially, that selective pressures in pre-industrial England favoured the genes and culture of the rich over the poor and criminal, that the former partly displaced the latter, and this partly explains why England became thriftier, more literate and less violent (*A Farewell to Alms*, 2008).)

On economics: Beinhocker's *Complexity Economics*, Mandelbrot's *The Misbehaviour of Financial Markets*, Vernon Smith's *Rationality*, Mackenzie's *An Engine, Not a Camera* (on financial models; an essay). This paper by Lo describes differences between the treatment of uncertainty in physics and economics. This memoir by Charles Ferguson gives a fascinating glimpse of how a Silicon Valley start-up really happens.

On 'complex systems' and the maths of prediction: this talk by Fields Medallist Terry Tao, Gell Mann's *The Quark and the Jaguar*, E.O. Wilson's *Consilience*, Barabasi's *Linked*. Mitchell's *Complexity, 2009* covers many of the inter-disciplinary subjects from computer science to the immune system (some maths is useful but not necessary to learn a lot from this book). Recently Nate Silver, who changed political coverage by using his blog to demonstrate mathematically robust political predictions, wrote *The Signal and the Noise* (2012) which is an excellent introduction to many ideas about modelling and prediction.

On war and international relations… Three classics: Thucydides' *The History of the Peloponnesian War;* Sun Tzu's *The Art of War*; Clausewitz's *On War*… Three modern works: van Creveld's *The Transformation of War*, Colin Gray's *Introduction to Strategic History*, John Robb's *Brave New War* (how civilian technological trends are changing modern war). One of Britain's most senior and respected civil servants wrote this paper after retiring on nuclear strategy. Anyone interested in terrorism and counter-terrorism must watch Pontecorvo's brilliant movie *The Battle of Algiers*, regularly found in terrorist safe houses. Perhaps the best modern subject for those interested in how political decisions are taken is Bismarck: Pflanze's *Biography*. On the difficulties of strategic thinking in the modern British state, Hoskyns' *Just in Time*. From inside a successful modern political campaign, Plouffe's *Audacity to Win*. To understand how science and technology has invaded political campaigns, read *The Victory Lab* (2012). Cialdini's *Influence* is interesting on the psychology of selling and marketing. To understand modern 'communications', read James Frayne's *Meet the People* (forthcoming). Almost the only good business book I have read on strategy and decisions is *Good Strategy, Bad Strategy* by Rumelt (2012). Gawande's *The Checklist Manifesto* is very interesting on how some fields (e.g airlines, medicine) have significantly improved performance while others (e.g education) have not, and the barriers to improvement. Tufte's *The Visual Display of Quantitative Information* and *Beautiful Evidence* are classics about the presentation of information.