



BANK OF ENGLAND

# Speech

---

## Tails of the unexpected

Paper by

Andrew G Haldane, Executive Director, Financial Stability and member of the Financial Policy Committee and Benjamin Nelson, Economist, Financial Stability

Given at “The Credit Crisis Five Years On: Unpacking the Crisis”, conference held at the University of Edinburgh Business School, 8-9 June

8 June 2012

The views are not necessarily those of the Bank of England or the Financial Policy Committee. I would like to thank Nick Chadwick-Williams, Iryna Kaminska, Priya Kothari, Dmitry Kuvshinov, Joseph Noss, Harshil Shah, Marco Spaltro, Peter Zimmerman and Gabriele Zinna for comments and contributions.

For almost a century, the world of economics and finance has been dominated by randomness. Much of modern economic theory describes behaviour by a random walk, whether financial behaviour such as asset prices (Cochrane (2001)) or economic behaviour such as consumption (Hall (1978)). Much of modern econometric theory is likewise underpinned by the assumption of randomness in variables and estimated error terms (Hayashi (2000)).

But as Nassim Taleb reminded us, it is possible to be *Fooled by Randomness* (Taleb (2001)). For Taleb, the origin of this mistake was the ubiquity in economics and finance of a particular way of describing the distribution of possible real world outcomes. For non-nerds, this distribution is often called the bell-curve. For nerds, it is the normal distribution. For nerds who like to show-off, the distribution is Gaussian.

The normal distribution provides a beguilingly simple description of the world. Outcomes lie symmetrically around the mean, with a probability that steadily decays. It is well-known that repeated games of chance deliver random outcomes in line with this distribution: tosses of a fair coin, sampling of coloured balls from a jam-jar, bets on a lottery number, games of paper/scissors/stone. Or have you been fooled by randomness?

In 2005, Takashi Hashiyama faced a dilemma. As CEO of Japanese electronics corporation Maspro Denkoh, he was selling the company's collection of Impressionist paintings, including pieces by Cézanne and van Gogh. But he was undecided between the two leading houses vying to host the auction, Christie's and Sotheby's. He left the decision to chance: the two houses would engage in a winner-takes-all game of paper/scissors/stone.

Recognising it as a game of chance, Sotheby's randomly played "paper". Christie's took a different tack. They employed two strategic game-theorists – the 11-year old twin daughters of their international director Nicholas Maclean. The girls played "scissors". This was no random choice. Knowing "stone" was the most obvious move, the girls expected their opponents to play "paper". "Scissors" earned Christie's millions of dollars in commission.

As the girls recognised, paper/scissors/stone is no game of chance. Played repeatedly, its outcomes are far from normal. That is why many hundreds of complex algorithms have been developed by nerds (who like to show off) over the past twenty years. They aim to capture regularities in strategic decision-making, just like the twins. It is why, since 2002, there has been an annual international world championship organised by the World Rock-Paper-Scissors Society.

The interactions which generate non-normalities in children's games repeat themselves in real world systems – natural, social, economic, financial. Where there is interaction, there is non-normality. But risks in real-world systems are no game. They can wreak havoc, from earthquakes and power outages, to

depressions and financial crises. Failing to recognise those tail events – being fooled by randomness – risks catastrophic policy error.

So is economics and finance being fooled by randomness? And if so, how did that happen? That requires a little history.

## **A Short History of Normality**

### *(a) Normality in Physical Systems*

Statistical normality had its origin in a set of physical experiments in the 17<sup>th</sup> century. Galileo discovered many of its fundamental properties when studying the measurement of distances to the stars. He found that random errors were inevitable in instrumental observations. But these errors exhibited a distinctive pattern, a statistical regularity: small errors were more likely than large and were symmetric around the true value.

This “reversion to the mean” was formalised in 1713 by Jacob Bernoulli based on a hypothetical experiment involving drawing coloured pebbles from a jam-jar.<sup>1</sup> The larger the number of drawings, the greater the chance that the observed number of coloured balls approximated its true average value. Although the game was simple, its implications were far-reaching. By simply taking repeat samplings, the workings of an uncertain and mysterious world could seemingly be uncovered.

The date when the normal distribution was formally introduced is known with precision. On November 12, 1733, the second edition of *The Doctrine of Chances* by Abraham de Moivre first plotted the bell curve. This was based on the probability distribution of another repeat game of chance – the number of “heads” which appeared when a fair coin was tossed. Such was the beauty of the distribution uncovered by this experiment, de Moivre “ascribed it to the Almighty”.<sup>2</sup>

Understanding of the normal curve was advanced in the early 19<sup>th</sup> century by two mathematicians, Carl Friedrich Gauss and Pierre Simon de Laplace. Gauss, a maths genius, was using the curvature of the Earth to improve the accuracy of geographic measurements in the Bavarian hills. He observed that the distribution of estimates varied widely, but tended to cluster around a mean with symmetry either side.

Gauss also forged a link between the normal curve and a particular approach to statistical inference – the “least squares” method. He pioneered the least squares method to locate the dwarf planet *Ceres*, minimising random errors from astronomical observations earlier provided by Italian monk Piazzi. This link between normality and least squares has continued to this day.

---

<sup>1</sup> Bernstein (1998).

<sup>2</sup> Bernstein (1998).

In 1810, in parallel to Gauss's work, Laplace provided a more theoretically sound justification for the use of the normal, achieved without the need for a jam-jar, coin toss or Bavarian hillside. He showed mathematically that the sum of a large number of mutually independent, identically distributed random variables is approximately normally distributed. Laplace's findings were the first appearance of what is now known as the central limit theorem.

To scientists seeking to explain the world, the attraction of the normal curve was obvious. It provided a statistical map of a physical world which otherwise appeared un-navigable. It suggested regularities in random real-world data. Moreover, these patterns could be fully described by two simple metrics – mean and variance. A statistical window on the world had been opened. To 19<sup>th</sup> century nerds, the Gaussian world was their oyster.

### *(b) Normality in Social Systems*

Not surprisingly, the normal distribution quickly caught on as it began to be applied in social and biological settings. Starting in the early 1830s, Belgian sociologist Adolphe Quetelet began to apply the normal distribution to the study of all manner of human attributes, physical and mental – from the chest size and height of Scottish and French soldiers, to murder and suicide rates, and to the first body mass index. Human attributes, like planetary attributes, seemed also to follow the bell curve's contours.

Quetelet was quick to cotton on to the potential. He has been branded the greatest regularity salesman of the 19<sup>th</sup> century (Hacking (1990)), developing the concept of *l'homme moyen* (the average man) to describe the social characteristics of the population at large. 150 years later, *l'homme moyen* was to reappear in economics as the representative agent.

It was not just de Moivre who was taken by the powers of the bell curve. For English statistician Francis Galton, the regularities of the normal curve held mesmeric qualities. It:

“... reigns with serenity and in complete self-effacement amidst the wildest confusion. The huger the mob...the more perfect its sway. It is the supreme law of Unreason. Whenever a large sample of chaotic elements are taken in hand...an unsuspected and most beautiful form of regularity proves to have been latent all along”.<sup>3</sup>

If the Greeks had known of the normal curve, Galton conjectured, they would have deified it. Not that all of Galton's work was exactly saintly. He advocated the use of data on personal characteristics for the dark purposes of eugenics. In his work *Hereditary Genius*, Galton included an estimate of the proportion of the population which could be classified as “eminent” based on their attributes – 1 in every 4,000.

---

<sup>3</sup> Bernstein (1998).

Galton's study of hereditary characteristics provided a human example of Bernoulli's reversion to the mean. Outlying human characteristics – height, weight, eyelash length, bad breath – were reduced over time as they passed down the generations. This Galton called “regression towards mediocrity in hereditary stature”. Hereditary regression is the foundation for the modern-day econometric concept of statistical regression.

Interestingly, up until this point Gauss's distribution did not have a name. Multiple independent appearances of the “normal” description then occurred at around the same time in three different countries (Peirce (1873), Galton (1877), Lexis (1877)).<sup>4</sup> C.S. Peirce drew smooth densities through the measurement errors made when detecting a sharp sound generated by dropping an object onto a wooden base and declared the distribution “normal”. By the end of the 19<sup>th</sup> century, the name had stuck.

This semantic shift was significant. The word “normal” acquired its contemporary meaning of “usual” towards the middle of the 19<sup>th</sup> century. It then quickly entered the popular lexicon. Statistical normality emerged as part of this etymological journey. From simple games of chance, normality found applications in everything from astronomy to public affairs to social choice. In the 18<sup>th</sup> century, normality had been formalised. In the 19<sup>th</sup> century, it was socialised. The normal distribution was so-named because it had become the new norm.

Up until the late 19<sup>th</sup> century, no statistical tests of normality had been developed. Having become an article of faith, it was deemed inappropriate to question the faith. As Hacking put it, “thanks to superstition, laziness, equivocation, befuddlement with tables of numbers, dreams of social control, and propaganda from utilitarians, the law of large numbers became a synthetic *a priori* truth”.<sup>5</sup> We were all Gaussians now.

### *(c) Normality in Economic and Financial Systems*

At about this time, the fledgling discipline of economics was emerging. Early models of the economic system developed by Classical economists were qualitative and deterministic. This followed the tradition in Newtonian physics of explaining the world using Classical deterministic laws. Jevons, Walras, Edgeworth and Pareto “transmuted the physics of energy into the social mechanics of utility” (Mirowski (1989)).

But in the early part of the 20<sup>th</sup> century, physics was in the throes of its own intellectual revolution. The emergence of quantum physics suggested that even simple systems had an irreducible random element. In physical systems, Classical determinism was steadily replaced by statistical laws. The natural world was suddenly ruled by randomness.

Economics followed in these footsteps, shifting from models of Classical determinism to statistical laws. The foundations for this shift were laid by Evgeny Slutsky (1927) and Ragnar Frisch (1933). They divided the

---

<sup>4</sup> Stigler (1999).

<sup>5</sup> Hacking (1990).

dynamics of the economy into two elements: an irregular random element or *impulse* and a regular systematic element or *propagation* mechanism. This impulse/propagation paradigm remains the centrepiece of macro-economics to this day.

The Frisch-Slutsky apparatus provided a conceptual framework within which empirical macro-economic models could commence their journey. In 1932, Alfred Cowles established the Cowles Commission for Research in Economics. The Cowles Foundation, as it became known in the 1950s, pioneered the development of large-scale macro-econometric models, including for policy analysis.

At the core of the macro-economic models developed by the Cowles Foundation were two features: least-squares methods for estimation of the economy's propagation mechanisms and normality of the random impulses to this system. Both would have been familiar to Gauss and Galton. These two elements were in principle separable – the least squares method does not in fact rely on any distributional assumption about errors. But in the years that followed, econometric estimation and normality became inseparable.

As evidence of that, tests for normality began to be applied to econometrically-estimated models in the 1970s. Tellingly, these tests were used as a diagnostic check on the adequacy of the model. A finding of abnormality in the errors, in economics as in psychiatry, was assumed to imply a failure of the model, a disease requiring treatment. As in the natural sciences in the 19<sup>th</sup> century, far from being a convenient statistical assumption, normality had become an article of faith. Normality had been socialised.

This line of thinking moved effortlessly from economics to finance. Harry Markowitz was a member of the Cowles Commission. In 1952, he wrote a paper which laid the foundations for modern portfolio theory (Markowitz (1952)). In line with his Cowles contemporaries, Markowitz assumed financial returns could be characterised by mean and variance alone – conveniently consistent with normality. That assumption was crucial, for from it followed Markowitz's mean-variance optimal portfolio rule.

At around the same time, Kenneth Arrow and Gerard Debreu (1954) were developing the first genuinely general equilibrium economic model. In this Arrow-Debreu world, future states of the world were assumed to have knowable probabilities. Agents' behaviour was also assumed to be known. The Arrow-Debreu model thereby allowed an explicit price to be put on risk, while ignoring uncertainty. Risky (Arrow) securities could now be priced with statistical precision. These contingent securities became the basic unit of today's asset pricing models.

In the period since, the models of Markowitz and Arrow/Debreu, with embedded assumptions of normality, have dominated asset-pricing in economics and finance. In economics, the Arrow/Debreu equilibrium model is the intellectual antecedent of today's real business cycle models, the dominant macro-economic

framework for the past 20 years (for example, Kiyotaki (2011)). Typically, these models have Gaussian-distributed impulses powering a Quetelet-inspired representative agent.

In finance, the dominant pricing models are built on Markowitz mean-variance foundations and the Arrow-Debreu principle of quantifiable risk. They, too, are typically underpinned by normality. For example, the feted Black and Scholes (1973) options-pricing formula, itself borrowed from statistical physics, is firmly rooted in normality. So too are off-the-shelf models of credit risk, such as Vasicek (2002). Whether by accident or design, finance theorists and practitioners had by the end of the 20<sup>th</sup> century evolved into fully paid-up members of the Gaussian sect.

### Assessing the Evidence

Now for an abnormal question: to what extent is normality actually a good statistical description of real-world behaviour? In particular, how well does it describe *systems* of behaviour, whether natural, social, economic or financial? Evidence against has been mounting for well over a century.

In the 1870s, the German statistician Wilhelm Lexis began to develop the first statistical tests for normality. Strikingly, the only series Lexis could find which closely matched the Gaussian distribution was birth rates. In 1929, E B Wilson and M M Hilferty re-examined using formal statistical techniques the datasets used by Peirce (Wilson and Hilferty (1929)). The distribution of the original data was found to be incompatible with the normal model. The natural world suddenly began to feel a little less normal.

The period since has seen mounting evidence of non-normality in a variety of real-world settings. In its place, natural and social scientists have often unearthed behaviour consistent with an alternative distribution, the so-called power law distribution. Mathematically, this takes the form:

$$P(X > x) \sim x^{-\alpha}$$

The probability that random variable  $X$  exceeds some level  $x$  is proportional to  $1/x^\alpha$ . In other words, the probability of large events decays polynomially with their size. This sounds innocuous enough. But in a Gaussian world, the probability of large events decays *exponentially* with their size, making large events increasingly rare at a rapid rate. Under power laws, these large events are much more likely.

This behaviour in the tail of the distribution makes power laws distinctive. Fat tails are a regularity. The speed of decay in the tail is governed by the parameter  $\alpha > 0$ . When  $\alpha$  falls, these tails fatten. And when  $\alpha$  takes a low value, many of the concepts familiar from a Gaussian world are turned on their head. Under the normal distribution means and variances are all that matter. For power laws with sufficiently fat tails, the

mean and variance may not even exist. Technically, power law-distributed data only have a well-defined mean when  $\alpha$  lies above unity and their variance only exists when  $\alpha$  exceeds two.

In consequence, Laplace's central limit theorem may not apply to power law-distributed variables. There can be no "regression to the mean" if the mean is ill-defined and the variance unbounded. Indeed, means and variances may then tell us rather little about the statistical future. As a window on the world, they are broken. With fat tails, the future is subject to large unpredictable lurches - what statisticians call kurtosis.

To bring this to life, consider the statistical properties of a set of natural, social and economic systems, beginning with the natural world. The left-hand panels of Charts 1-3 show estimated probability densities for solar flares, earthquakes and rainfall.<sup>6</sup> The right-hand side panel shows a transformation of these data. Conveniently, this allows a simple comparison of the data with the normal curve which appears as a straight line. So any deviation from the straight line at the extremities signals a non-normal fat tail.<sup>7</sup>

From Charts 1-3, all three natural phenomena clearly exhibit a fat upper tail: incidences of exceptionally large solar flares, earthquakes and rainfall occur more frequently than would be implied by the normal distribution. Table 1 shows a statistical measure of kurtosis. Under normality, this statistic would equal 3. For these natural systems, kurtosis measures are larger, sometimes far larger, than their normal equivalents.

This evidence strongly suggests that the laws of physics can generate highly non-normal system-wide behaviour. Even in systems without human behaviour, large-scale unexpected events – fat tails – appear to be a prominent feature of the environmental landscape. Assuming the physical world is normal would lead to a massive under-estimation of natural catastrophe risk.

Moving from natural to social systems, evidence of fat-tails is equally strong. Charts 4-6 look at three social systems: the frequency with which different words appear in the novel *Moby Dick*; the frequency of citations of pieces of scientific research; and the population size of US cities.<sup>8</sup> While seemingly disparate, each is in its own way shaped by a common human factor – social interaction.

Graphically, it is clear that all three of these social systems have a large upper tail. Indeed, they are all typically found to be power law-distributed. For city size, this distribution goes by the name Zipf's Law (as first noted by Auerbach (1913)). It has a striking pattern: the largest city is twice the size of the second-largest, three times the size of the third and so on. A comparably self-similar pattern is also found in the distribution of names, wealth, words, wars and book sales, among many other things (Gabaix (2009)).

---

<sup>6</sup> The Data Annex provides a definition of these series.

<sup>7</sup> More precisely, it signals a fat tail if the observed distribution lies to the left of the straight line at the bottom of the chart and to the right of the straight line at the top of the chart. This means that extreme events have a larger probability than that suggested by the normal distribution.

<sup>8</sup> The data source is Newman (2005).

The kurtosis statistics in Table 1 indicate that fat-tailedness is considerable in both human and natural systems. For example, for earthquakes it is around 3000 times greater than normal, for city sizes 2000 times greater. As a hybrid, we might expect non-normalities in economic and financial systems too. To assess that, Charts 7-10 plot four measures drawn from those systems: GDP, rice prices, bank credit and equity prices. They are drawn from a range of countries and measured over a lengthy sample – credit and GDP over more than a century, equity prices over more than three centuries and rice prices over a millennium.<sup>9</sup>

In all four of these series, there is visual evidence of fat tails. The statistical measures of kurtosis in Table 1 confirm that. Tails appear to be fatter in the financial (equity returns and bank credit) than macro-economic series (GDP). But for both, the fatness of the tails is significant.

To bring to life the implications of these fat tails, consider an insurance contract designed to guard against catastrophes in the tail of the distribution of outcomes. In particular, assume that this insurance contract only pays out if outcomes are more than four standard deviations above their mean value in any one year (or below for the economic series). Under normality, payouts would be expected very rarely.<sup>10</sup>

Now consider the actuarially-fair value of the insurance premium for this contract. This can be calculated under two different assumptions: first, assuming normality of the distribution of outcomes, and second using the observed, fat-tailed distribution of outcomes. The ratio of the difference between the resulting insurance premia is shown in Table 2.

These differences are enormous. For economic and financial series, they are typically multiples of 100 or more. This suggests the assumption of normality would result in massive under-pricing of catastrophe insurance risk, by two orders of magnitude. This mis-pricing is as or more acute when insuring against economic catastrophes (such as output crashes) as natural catastrophes (such as earthquakes).

Put differently, consider the implied probabilities of a three-sigma fall in GDP or equity prices. Assuming normality, catastrophe risk on this scale would be expected to occur approximately once every 800 years for GDP and once every 64 years for equities. In reality, for GDP it appears to occur roughly once every century, for equities once every 8 years.

## Explaining Fat Tails

So what explains these fat-tailed outcomes in natural and social systems? One factor holds the key - interactions. The central limit theorem is predicated on the assumption of independence of observations. In complex systems, natural and social, that assumption is almost certain to be violated. Systems are systems

---

<sup>9</sup> The Data Annex provides a definition of these series.

<sup>10</sup> Around once every 31,000 years.

precisely because they are *interdependent*. In a nutshell, that is why so few systems behave normally in the statistical sense. This interaction can take a variety of forms.

*(a) Non-Linear Dynamics*

As one example from the physical world, consider weather systems. The first meteorological offices were established in the mid-19<sup>th</sup> century to help anticipate extreme weather events which might interrupt shipping trade. Up until the late 19<sup>th</sup> century, these offices used isobaric charts as the basis for forecasting (Davis 1984). Physicists often viewed meteorologists with scepticism. They were the phrenologists of science, reading the isobaric bumps.

That changed at the start of the 20<sup>th</sup> century. In 1904, Vilhelm Bjerknes established a set of seven equations with seven unknowns governing the evolution of the atmosphere. This provided the theoretical foundation for meteorology. This system itself was a complex one, with non-linearities which arose from interactions among atmospheric convections. But if initial conditions for the system were properly established, this system could be solved using numerical methods to yield a deterministic path for future weather.

This, it transpired, was a big if. In 1963, American meteorologist Edward Lorenz was simulating runs of weather predictions on his computer. Returning from his coffee break, he discovered the new run looked completely different than the old one. He traced the difference to tiny rounding errors in the initial conditions. From this he concluded that non-linear dynamic systems, such as weather systems, exhibited an acute sensitivity to initial conditions. Chaos was born (Gleick (1987)).

Lorenz himself used this chaotic finding to reach a rather gloomy conclusion: forecasts of weather systems beyond a horizon of around two weeks were essentially worthless. And the gloom soon spread. Following Lorenz's chance finding, chaotic dynamics were unearthed throughout the natural and human world, ranging from gravitational forces to population growth.

These chaotic dynamics carry important implications for system-wide behaviour. Small changes in the system can lead to radically different real-world outcomes. Tipping point or, in Lorenz's words, butterfly effects proliferate. Unlike in much of economics and finance, equilibria are often neither singular nor stable. Reversion to the mean is a poor guide to the future, often because there may be no such thing as a fixed mean. Fat tails sprout in every direction.

Economics and finance is not a complete stranger to models of non-linear dynamics. During the 1980s, they were used to study the business cycle (Grandmont (1985)), optimal growth (Day (1983)) and consumer choice (Benhabib and Day (1981)). But the emergence of real business cycle theory, with its linear decision rules and representative agent, has snuffed-out this effort during much of the past two decades.

The past five years have seen plenty of evidence of Lorenz-like non-linearities in the economy and the financial system. Chaos, metaphorical and mathematical, has been to the fore. It is not difficult to identify some of the structural features which may have contributed to these chaotic dynamics. Leverage would be one. The accumulation of leverage was a key feature of the pre-crisis boom and subsequent bust. Leverage generates highly non-linear system-wide responses to changes in income and net worth (Thurner et al (2010)), the like of which would have been familiar to Lorenz.

*(b) Self-Organised Criticality*

In weather systems, chaotic dynamics and tipping points are embedded in the physical laws controlling the system. In others, the system automatically organises itself into a fragile state, perennially perched on a knife-edge. These are systems of so-called self-organised criticality (Bak (1996)).

The classic example is the sand pile. Imagine constructing a sand pile, grain by grain. Each grain is independent and identically distributed, Laplace-style. But the distribution of this system does not conform to the central limit theorem. As grains are added, mini-cascades begin. At some point, these avalanches offset the effects of adding extra grains. The sand pile has reached its self-organised critical state.

Another example from the natural world would be forest fires. Imagine a plot of land in which trees are randomly planted over time, which is subject to occasional lightning strikes that could cause a fire. When a fire starts, it spreads to adjoining trees. Over time, the forest will organise itself into a critical state where, most of the time, fires are small and contained. But on occasions, a random lightning strike can cause the whole forest to be lost. The system then resumes its self-organising march back towards the critical state.

These self-organised criticalities have been found in human systems too. Left to its own devices, traffic tends to converge on a critical speed at which the risk of a log-jam or pile-up are dramatically increased. As in sand piles and forest fires, traffic dynamics are subject to phase transitions. Adding one extra grain of sand, one well-directed lightning bolt, one speeding car, can push the system from its critical state into fat-tailed catastrophe.

Self-organised criticality has also found itself into some of the darker corners of economics and finance. Bak, Chen, Scheinkman and Woodford (1993) develop a model in which shocks to an individual firm's production are independent and identically distributed, like individual grains of sand. But the supply chain among firms self-organises to generate economy-wide fluctuations, like sand pile cascades. In finance, Benoit Mandelbrot (1963) found that equity prices exhibit behaviour consistent with self-organised criticality.

The build-up to the financial crisis provided another example of self-organised criticality. A competitive search for yield among financial firms caused them to increase risk to hit return targets. The sand pile of

returns grew ever-higher as the system self-organised on a high-risk state. Having reached this critical state, a small addition of risk – a single sub-prime grain of sand – was sufficient to cause the whole edifice to collapse in an uncontrolled cascade (Haldane (2012a)).

*(c) Preferential Attachment*

Neither non-linear dynamics nor self-organised criticality rely on human behaviour or social interaction. Introducing human behaviour is likely to make for stronger interactions within the system, further fattening the tail. Social networks, unlike physical or biological networks, are in many ways defined by these interactions. Without them, the network simply would not exist.

Social networks, be it school classrooms, churches, pubs or the world wide web, have been extensively studied (Jackson (2010)). As a network, they exhibit some common topological features. For example, most have a large number of poorly-connected agents and a relatively small number of highly-connected agents. Why so? One explanation is so-called preferential attachment.

Imagine a network of nodes initially connected in a random way. Now assume a new node is added which links to existing nodes with the highest degree of connectedness. There is preferential attachment. Intuitively, it is easy to see why this might be. Popularity is contagious and self-fulfilling. The resulting network will be characterised by a high degree of connectivity for a few central nodes. It will be power law-distributed (Barabasi and Albert (1999)).

Nowhere are the dynamics of network formation clearer than in the world wide web. People are much more likely to connect to a blogger or Tweeter with a large number of followers than a small number. Why? Because popularity can be a signal of quality – the Stephen Fry phenomenon. Or because, even if quality is low, popularity can be self-fulfilling – the Kim Kardashian phenomenon.

The same basic dynamic operates across most, if not all, social and socio-economic networks. Preferential attachment explains the distribution of web-links, academic citations and Facebook friends (Barabasi and Albert (1999)). It explains the distribution of city sizes (Zipf's Law). It also explains the outcomes from strategic games, for both adults (nuclear proliferation) and children (paper/scissors/stone). All are well-known to be power law distributed.

These types of preferential attachment have a history in economics too. Keynes viewed the process of forming expectations as more beauty pageant than super-computer (Keynes (1936)). Agents form their guess not on an objective evaluation of quality (Stephen Fry) but according to whom they think others might like (Kim Kardashian). Stephen Fry has just over 4 million Twitter followers. Kim Kardashian has 15 million.

This sheep-like logic makes for multiple equilibria, as expectations herd themselves into different pens. Some of these equilibria may be bad ones. The classic example in finance is the Diamond and Dybvig (1983) bank run. If depositors believe others will run, so will they. Financial unpopularity then becomes infectious. People queuing outside Northern Rock in 2007 were behaving identically to people following Kim Kardashian on Twitter in 2012. Both made for a highly sub-optimal equilibrium.

Co-ordination games such as these appear to operate in a wide variety of settings. The links can be physical, evolutionary, financial, social or expectational (Haldane (2012b)). Each gives rise to highly non-linear system dynamics, due to butterfly effects and phase transitions as expectations shift. The financial crisis has been riddled with examples of these sorts of behaviour, from Lehman Brothers' failure to the euro-zone crisis. In each case, expectational chains of panic and fear delivered bad equilibria.

#### *(d) Highly-Optimised Tolerance*

Some systems organise themselves into a critical state by themselves. Other are nudged into doing so by the hand of man. In other words, in some systems criticality may be man-made. These are often described as systems of highly-optimised tolerance (Carlson and Doyle (2002)).

Consider the forest fire model. But now introduce a forester who has some concern for the expected yield of the forest, proxied by the number of planted trees. The forester now faces a trade-off. A more densely-filled forest makes for superior expected yield. But it also exposes the forest to systemic fire risk, which could wipe out large areas of trees and damage future yield. How to strike the right balance?

The optimal response by the forester is to build in fire breaks. These should be larger in number in areas where lightning strikes are more frequent. In areas where they are rare, forestation can afford to be denser. This arrangement maximises expected yield. But it will also result in occasional systemic forest fires. Moreover, if the forester miscalculates the probability of lightning strikes, the system might be sub-optimally susceptible to catastrophic failure. Either way, the result will be a fat-tailed distribution of forest fire size.

These outcomes, where critical states are generated by human intervention, were also a feature of the financial crisis. Pre-crisis, regulators set capital ratios based on an assessment of the riskiness of banks' assets. These assessments were imperfect. With hindsight, assets that were underweighted in risk terms (trading book assets, sovereign debt) encouraged banks to invest in them. The financial sector organised itself on a critical, high-risk state due to well-meaning, but imperfect, regulation.

Taken together, the picture these models paint is a vivid one. It is not difficult to imagine the economic and financial system exhibiting some, perhaps all, of these features – non-linearity, criticality, contagion. This is particularly so during crises. Where interactions are present, non-normalities are never far behind. Indeed,

to the extent that financial and economic integration is strengthening these bonds, we might anticipate systems becoming more chaotic, more non-linear and fatter-tailed in the period ahead.

## **Where Next?**

Given that, what can be done to better recognise and manage non-normalities? Fat tails contain important lessons for economics and finance. They also contain some important lessons for economic and financial policymakers.

### *(a) Non-Normality in Economics and Finance*

Like de Moivre in the 18<sup>th</sup> century and Galton in the 19<sup>th</sup>, the economics profession has for much of the 20<sup>th</sup> century been bewitched by normality. Real business cycle theory in economics and efficient markets theory in finance bear the tell-tale signs of this intellectual infatuation. So too does much of econometrics. All three are rooted in a Frisch/Slutsky meets Arrow/Debreu framework, with normal impulses acting on linear propagation rules. Predictably, this generates near-Gaussian outcomes for macro-economic variables.

Over the past five years, the real world has behaved in ways which make a monkey of these theories. In the face of shocks, some of them modest, the economic and financial world has often responded in highly irregular and discontinuous ways. Tipping points and phase transitions have been the name of the game. The disconnect between theory and reality has been stark. Economics and finance, like Sotheby's, may have been fooled by randomness.

Taking a step forward may require economics and finance to first take a step back. In 1921, Frank Knight drew an important distinction between risk on the one hand and uncertainty on the other (Knight (1921)). Risk arises when the statistical distribution of the future can be calculated or is known. Uncertainty arises when this distribution is incalculable, perhaps unknown.

Many of the biggest intellectual figures in 20<sup>th</sup> century economics took this distinction seriously. Indeed, they placed uncertainty centre-stage in their policy prescriptions. Keynes in the 1930s, Hayek in the 1950s and Friedman in the 1960s all emphasised the role of uncertainty, as distinct from risk, when it came to understanding economic systems. Hayek criticised economics in general, and economic policymakers in particular, for labouring under a "pretence of knowledge" (Hayek (1974)).

Yet it is risk, rather than uncertainty, that has dominated the economics profession for much of the past 50 years. By assuming future states of the world were drawn from known distributions, Arrow and Debreu enabled risk to be priced with statistical precision and uncertainty to be conveniently side-stepped.

Uncertainty was, quite literally, ruled out of the equation. But if economic and financial systems operate on the border between order and disorder, ignoring uncertainty is deeply unrealistic.

Uncertainty profoundly alters the way systems behave. Take asset pricing. Under uncertainty rather than risk, asset prices are no longer defined by a single price. Instead their equilibrium price is defined by a range (Epstein and Wang (1994)). Prices systematically differ from their fundamental values. If uncertainties escalate, they experience phase shifts. Caballero and Krishnamurthy (2008) have examined the implications of financial uncertainty for systemic risk.

In response to the crisis, there has been a groundswell of recent interest in modelling economic and financial systems as complex, adaptive networks. For many years, work on agent-based modelling and complex systems has been a niche part of the economics and finance profession. The crisis has given these models a new lease of life in helping explain the discontinuities evident over recent years (for example, Kirman (2011), Haldane and May (2011)).

In these frameworks, many of the core features of existing models need to be abandoned. Quetelet's *l'homme moyen* is replaced by interactions among non-representative agents, Twitter-style. The single, stationary equilibrium gives way to Lorenz-like multiple, non-stationary equilibria. Frisch-Slutsky linearity is usurped by sand pile-style tipping points. In economics and finance, these system-wide behaviours may today be close to the analytical rule.

These types of systems are far from alien territory to physicists, sociologists, ecologists and the like. Were they to enter the mainstream economics and finance, they would give it a realistic prospect of generating the sorts of fat-tailed non-normalities present in the real world. Doing so will require a fairly fundamental re-think of the foundations of modern-day economics, finance and econometrics.

#### *(b) Non-Normality and Risk Management*

The risk management tools used by financial institutions have, in many ways, a greater distance still to travel. As an example of that, consider the Value-at-Risk (or VaR) techniques widely used by banks for capturing and managing portfolio risk.

VaR is a statistical measure of risk developed by JP Morgan in the 1990s. It measures the maximum loss from a given portfolio position at a certain confidence level over a certain period of time. For example, if a bank's 10-day 99% VaR is \$3 million, there is considered to be only a 1% chance that losses will exceed \$3 million over 10 days. In this role, VaR can be used to help set risk limits for traders' portfolios. It can also be used to set regulatory capital standards for these portfolios.

The simplicity of VaR has led to its ubiquitous use in finance (Jorion (2006)). But VaR suffers a fatal flaw as a risk management and regulatory measure: it is essentially silent about risks in the tail beyond the confidence interval. For example, even if a trader's 99% VaR-based limit is \$10 million, there is nothing to stop them constructing a portfolio which delivers a 1% chance of a \$1 billion loss. VaR would be blind to that risk and regulatory capital requirements seriously understated.

Worse still, the fatter the tails of the risk distribution, the more misleading VaR-based risk measures will be. Consider holding a portfolio of world equities and, based on data from 1693 to 2011, calculate the VaR. The 99% VaR assuming the data are normal gives a loss of \$6 trillion at today's prices. Using the actual data raises the estimated VaR by one third to \$7.8 trillion. Finally, calculating the risk conditional on being in the 1% tail of the distribution gives a loss of \$18.4 trillion. Simple VaR underestimates risk by a factor of 1.5 and 3.

This example is far from hypothetical. The inadequacies of VaR were amply illustrated during the crisis. It could be argued that these crisis lessons have been learned. That is far from clear. In May this year, *Risk* magazine asked risk managers whether VaR should be scrapped. A majority said no. And as a measure of tail risk, VaR continues to surprise on the downside.

On 10 May, JP Morgan announced losses totalling \$2 billion on a portfolio of corporate credit exposures. This took the world, and JP Morgan's management, by surprise. After all, prior to the announcement the 95% VaR on this portfolio in the first quarter of 2012 had been a mere \$67 million. This VaR measure was revised upwards to \$129 million on the day of the announcement. Whether this proves a more accurate measure of tail risk on this still-open position remains to be seen.

These asset pricing and risk management problems do not begin and end with VaR. Many financial instruments have option-like payoff schedules. Since the early 1970s, these risks have typically been priced using the Black-Scholes options-pricing formula. But that model assumes normality of returns. If financial returns are instead power law-distributed, Black-Scholes – if taken at face value – can lead to a material mis-pricing of risk.

One example of that is a well-known puzzle in finance known as the "volatility smile". The volatilities implied by the Black-Scholes formula typically exhibit a "smile" pattern. In other words, implied volatilities are greater at the extremities – for example, for deeply out-of-the-money options. One explanation of the puzzle is that the underlying distribution of returns is fat-tailed. That would tend to boost the value of deeply out-of-the-money options relative to the Black-Scholes formula and thus raise Black-Scholes-implied volatilities.

Imagine pricing a call option on rice prices using data over the past 1000 years. Assume the price today is \$100 and the option expiring in a year's time is at a deeply out-of-the money strike price three standard deviations above the mean. Assuming a normal distribution for price changes, this option would cost one cent. If instead the actual distribution of rice prices were used, its fair-value price would be 37 cents. The normality embedded in Black-Scholes would generate an enormous mis-pricing of catastrophe risk.

Turning from market to credit risk, the off-the-shelf model is Vasicek (1991). In the basic version, the correlation of each asset in a bank's portfolio with a common risk factor is estimated. The Vasicek model is at the core of many large banks' risk management frameworks. For example, it is typically used to determine their capital requirements under the internal ratings-based (IRB) Basel framework (BCBS (2005)). But in generating a distribution of expected portfolio losses, standard applications of the Vasicek model assume that the underlying risk factors, and hence portfolio losses, are normally distributed.

Suppose the common factor in the Vasicek model is GDP growth and this is assumed to be normally distributed. Under a baseline calibration using the foundation IRB approach of Basel II, and assuming a tail outcome for the common factor that occurs with probability 0.1%, the baseline capital buffer held to cover unexpected losses would be around 3%.<sup>11</sup>

Now ask what happens if the actual, fat-tailed distribution of GDP over the past three centuries is used. Under the baseline calibration, this raises the required capital buffer four-fold to around 12%. Tarashev and Zhu (2008) undertake a formal study of the degree of credit risk mis-pricing under the Vasicek model. They find that capital requirements could be between 20% and 85% higher assuming fat-tailed distributions.

Another favoured empirical approach to modelling credit risk, in particular correlations across assets in structured credit products, is the so-called copula approach (Noss (2010)). The copula describes the co-dependence between asset values and is typically assumed to be Gaussian. If the joint distribution between asset values is instead fat-tailed, the copula will systematically mis-price risk in structured credit instruments. So it was during the crisis, as small differences in correlation made for dramatic adjustments in price.

Taken together, non-normality suggests the need for a fairly fundamental rethink of the core risk management tools currently used by many financial firms. That includes, importantly, models used to set regulatory capital requirements. Even post-crisis, too many of these models remain at risk of being misled by normality, fooled by randomness. That was a key fault-line during the crisis and, as recent experience attests, remains a key fault-line today.

---

<sup>11</sup> Assuming loss given default of 45%, probability of default of 1% and asset correlation of 10%.

*(c) Non-Normality and Systemic Risk*

Some of the tail risks facing financial firms are genuinely difficult to calibrate accurately. That is because they are created endogenously within the system as a result of the behaviour of other participants (Danielsson et al (2009)). Because those behaviours are unobservable, so too are the tail risks facing individual banks. That is a potentially serious risk management gap.

This gap can most obviously be filled by some systemic oversight agency, able to monitor and potentially model the moving pieces of the financial system. Pre-crisis, there were few, if any, such systemic regulatory bodies in place. But over the past few years, a number have emerged charged with just this task - the Financial System Oversight Council (FSOC) in the US, the European Systemic Risk Board (ESRB) in Europe and the Financial Policy Committee (FPC) in the UK, to name three.

One useful role these bodies can play is to provide a guide to the contours of systemic risk – a systemic risk map. This map could provide a basis for risk management planning by individual financial firms. As in weather forecasting, the systemic risk regulator could provide early risk warnings to enable defensive actions to be taken. Indeed, the evolution of weather forecasting may provide useful lessons on the directions finance might take – and some grounds for optimism.

After the Second World War, meteorologists were among the earliest adopters of computing technology. The processing power of computers has certainly been key to advances in weather forecasting. The main reason for Lorenz's pessimism about forecasting was the problem of collecting and processing data. Modern computers have stretched significantly those data limits. The UK Met Office's computer can handle 100 trillion calculations per second and uses hundreds of thousands of global observations each day.

The results have been striking. Errors in weather-forecasting have been in secular decline. Today's four-day forecasts are as accurate as one-day forecasts 30 years ago. Weather predictions have predictive value well beyond a two-week horizon. Indeed, predictions of climatic patterns are now made out to a horizon of 100 years.<sup>12</sup> Lorenz's pessimism was misplaced.

Finance could usefully follow in these footsteps. There are already steps underway internationally to widen and deepen the array of financial data available to systemic risk regulators, filling gaps in the global network map. As in weather forecasting, this would help improve estimates of the initial conditions of the financial system. And as in weather forecasting, it is important these data are captured in a common financial language to enable genuinely global maps to be drawn (Ali, Haldane and Nahai-Williamson (2012)).

---

<sup>12</sup> For example, see the *Fourth Assessment Report* of the Intergovernmental Panel on Climate Change (IPCC, 2007).

These data then need to be brought together using a set of behavioural models of the economic and financial system. Economics does not have the benefit of meteorologists' well-defined physical laws. But by combining empirically-motivated behavioural rules of thumb, and balance sheets constraints, it should be possible to begin constructing fledgling models of systemic risk.<sup>13</sup>

At the Bank of England, work has been underway on such a model since well before the crisis. The Risk Assessment Model for Systemic Institutions (RAMSI) combines macro-economic and financial network data to enable simulations and stress-tests of the UK financial system (Aikman et al (2009)). It generates system-wide risk metrics and distributions that could be used to warn about impending threats. Tellingly, the distributions of outcomes in the financial system tend to have highly non-normal, fat tails.

Navigating the contours of systemic risk is one thing. Re-profiling these contours is quite another. International regulators have only recently begun the task of calibrating regulatory rules with an eye to systemic, as distinct from institution-specific, risk. There is a considerable distance to travel. Regulatory rules of the past sought to reflect risk. Regulatory rules of the future will need to seek to reflect uncertainty.

That calls for a quite different, and sometimes seemingly perverse, approach. Under uncertainty, many of our intuitive regulatory rules of thumb are turned on their head: slower can be faster, less can be more, slack can be tight. Instinctively a complex financial system seems likely to require complex control rules. And under risk, that intuition is roughly right. Under uncertainty, however, it is precisely wrong. Then, the optimal control rule is typically a simple one (DeMiguel et al (2009)). Less is more (Gigerenzer and Brighton (2008)).

The reason less can be more is that complex rules are less robust to mistakes in specification. They are inherently fragile. Harry Markowitz's mean-variance optimal portfolio model has informed millions of investment decisions over the past 50 years – but not, interestingly, his own. In retirement, Markowitz instead used a much simpler equally-weighted asset approach. This, Markowitz believed, was a more robust way of navigating the fat-tailed uncertainties of investment returns (Benartzi and Thaler (2001)).

Regulation has begun to learn some of these lessons. The mainstay of regulation for the past 30 years has been more complex estimates of banks' capital ratios. These are prone to problems of highly-optimised tolerance. In part reflecting that, regulators will in future require banks to abide by a far simpler backstop measure of the leverage ratio. Like Markowitz's retirement portfolio, this equally-weights the assets in a bank's portfolio. Like that portfolio, it too will hopefully be more robust to fat-tailed uncertainties.

A second type of simple, yet robust, regulatory rule is to impose structural safeguards on worst-case outcomes. Technically, this goes by the name of a "minimax" strategy (Hansen and Sargent (2011)). The firebreaks introduced into some physical systems can be thought to be playing just this role. They provide a

---

<sup>13</sup> A recent example would be the calibrated, agent-based model of Geanakoplos et al (2012).

fail-safe against the risk of critical states emerging in complex systems, either in a self-organised manner or because of man-made intervention.

These firebreak-type approaches are beginning to find their way into the language and practice of regulation. The Volcker rule in the US and the Vickers proposals in the UK are very much in that spirit. Structural separation solutions are still looked upon with a mixture of scepticism and horror by many in the regulatory and risk management community. Certainly, on the face of it they appear rather crude as a risk management device.

Under uncertainty, however, that is precisely the point. In a complex, uncertain environment, the only fail-safe way of protecting against systemic collapse is to act on the structure of the overall system, rather than the behaviour of each individual within it. In averting avalanche risk, it makes no sense to set limits on the size and shape of every grain of sand. It makes considerable sense to assess, and perhaps reshape, the structure of the sand pile itself, however.

Finally, in an uncertain world, fine-tuned policy responses can sometimes come at a potentially considerable cost. Complex intervention rules may simply add to existing uncertainties in the system. This is in many ways an old Hayekian lesson about the pretence of knowledge, combined with an old Friedman lesson about the avoidance of policy harm. It has relevance to the (complex, fine-tuned) regulatory environment which has emerged over the past few years.

And the argument can be taken one step further. Attempts to fine-tune risk control may add to the probability of fat-tailed catastrophes. Constraining small bumps in the road may make a system, in particular a social system, more prone to systemic collapse. Why? Because if instead of being released in small bursts pressures are constrained and accumulate beneath the surface, they risk an eventual volcanic eruption. Taleb and Blyth (2011) apply this line of logic to explain the build-up and after-effects of the Arab Spring.

## **Conclusion**

Normality has been an accepted wisdom in economics and finance for a century or more. Yet in real-world systems, nothing could be less normal than normality. Tails should not be unexpected, for they are the rule. As the world becomes increasingly integrated – financially, economically, socially – interactions among the moving parts may make for potentially fatter tails. Catastrophe risk may be on the rise.

If public policy treats economic and financial systems as though they behave like a lottery – random, normal – then public policy risks itself becoming a lottery. Preventing public policy catastrophe requires that we better understand and plot the contours of systemic risk, fat tails and all. It also means putting in place

robust fail-safes to stop chaos emerging, the sand pile collapsing, the forest fire spreading. Until then, normal service is unlikely to resume.

## References

- Aikman, D, Alessandri, P, Eklund, B, Gai, P, Kapadia, S, Martin, E, Mora, N, Sterne, G, Willison, M (2009)**, "Funding liquidity risk in a quantitative model of systemic stability", *Bank of England Working Paper No. 372*.
- Ali, R D, Haldane, A G and Nahai-Williamson, P (2012)**, "Towards a common financial language", available at <http://www.bankofengland.co.uk/publications/Documents/speeches/2012/speech552.pdf>.
- Auerbach, F (1913)**, "Das Gesetz der Bevölkerungskonzentration", *Petermanns Geogr. Mitt.* 59,74–76.
- Arrow, K J and Debreu, G (1954)**, "Existence of an Equilibrium for a Competitive Economy", *Econometrica*, Vol. 22, No. 3, pp. 265-290.
- Bak, P, Chen, K, Scheinkman, J and Woodford, M (1993)**, "Aggregate fluctuations from independent sectoral shocks: self-organized criticality in a model of production and inventory dynamics", *Ricerche Economiche*.
- Bak, P (1996)**, "How nature works: the science of self-organized criticality", *Copernicus*.
- Barabasi, A L and Albert, R (1999)**, "Emergence of scaling in random networks", *Science*, 286(5439), pp. 509-12.
- Basel Committee on Banking Supervision (2005)**, "An Explanatory Note on the Basel II IRB Risk Weight Functions", *Bank for International Settlements*.
- Benartzi, S and Thaler, R (2001)**, "Naïve diversification strategies in defined contribution saving plans", *American Economic Review*, Vol 91, 1, pp.79-98.
- Benhabib, J and Day, R H (1981)**, "Rational Choice and Erratic Behaviour", *Review of Economic Studies* 48 (3), pp. 459-471.
- Bernstein, P (1998)**, "Against the Gods: the remarkable story of risk", *John Wiley & Sons*.
- Black, F and Scholes, M (1973)**, "The Pricing of Options and Corporate Liabilities", *Journal of Political Economy*, Vol 81 (3), pp. 637-654.
- Caballero, R J and Krishnamurthy, A (2008)**, "Collective Risk Management in a Flight to Quality Episode", *The Journal of Finance*, vol 63 (5), pp. 2195-2230.
- Carlson, J M and Doyle, J (2002)**, "Complexity and robustness", *PNAS* 99(1), pp. 2538-45.
- Cochrane, J H (2001)**, "Asset Pricing", *Princeton University Press*.
- Danielsson, J, Shin, H S and Zigrand, J-P (2009)**, "Risk Appetite and Endogenous Risk", available at [http://papers.ssrn.com/sol3/papers.cfm?abstract\\_id=1360866](http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1360866).
- Davis, J L (1984)**, "Weather Forecasting and the Development of Meteorological Theory at the Paris Observatory, 1853-1878", *Annals of Science*, 41(4) pp.359-382.
- Day, R H (1983)**, "The Emergence of Chaos from Classical Economic Growth", *Quarterly Journal of Economics*, 98, pp. 201-213.
- DeMiguel, V, Garlappi, L and Uppal, R (2009)**, "Optimal Versus Naive Diversification: How Inefficient is the 1/N Portfolio Strategy?", *The Review of Financial Studies*, 22(5), pp. 1915-1953.

- Diamond, D W and Dybvig, P H (1983)**, "Bank runs, deposit insurance, and liquidity", *Journal of Political Economy* 91(3), pp. 401–19.
- Epstein, L G and Wang, T (1994)**, "Intertemporal asset pricing under Knightian uncertainty", *Econometrica*, 62(3), pp. 283-322.
- Frisch, R (1933)**, "Propagation Problems and Impulse Problems in Dynamic Economics", in *Economic Essays in Honour of Gustav Cassel*. Eds.: London: Allen and Unwin, pp.171-205.
- Gabaix, X (2009)**, "Power laws in economics and finance", *Annual Review of Economics*, 1 pp. 255-93.
- Galton, F (1877)**, "Typical laws of heredity", *Nature*,15, pp. 492-553.
- Geanakoplos, J, Axtell, R, Farmer, D J, Howitt, P, Conlee, B, Goldstein, J, Hendrey, M, Palmer, N M, Yang, C-Y (2012)**, "Getting at systemic risk via an agent-based model of the housing market", *American Economic Review*, 102(3): 53–58, May.
- Gigerenzer, G and Brighton, H (2008)**, "Homo Heuristicus: Why Biased Minds Make Better Inferences", *Topics in Cognitive Science* (1), pp. 107-143.
- Gleick, J (1987)**, "Chaos: Making a New Science", *Viking*.
- Grandmont (1985)**, "On Endogenous Competitive Business Cycles", *Econometrica*, 53, pp. 995-1045.
- Hacking, I (1990)**, "The Taming of Chance", *Cambridge University Press*.
- Haldane, A G (2012a)**, "On counterparty risk", available at <http://www.bankofengland.co.uk/publications/Documents/speeches/2012/speech571.pdf>.
- Haldane, A G (2012b)**, "Financial arms races", available at <http://www.bankofengland.co.uk/publications/Documents/speeches/2012/speech565.pdf>.
- Haldane, A G and May, R M (2011)**, "Systemic risk in banking ecosystems", *Nature* (469), 351-355.
- Hall, R E (1978)**, "Stochastic implications of the life cycle-permanent income hypothesis", *Journal of Political Economy* 86 (6), pp. 971-87.
- Hansen, L P and Sargent, T (2011)**, "Wanting robustness in macroeconomics", in Friedman, B. M. and Woodford, M. (eds) *Handbook of Monetary Economics*, Vol 3B, North-Holland.
- Hayashi, F (2000)**, "Econometrics", *Princeton University Press*.
- Hayek, F A (1974)**, "The Pretence of Knowledge", *Nobel Memorial Prize Lecture*.
- IPCC (2007)**, "Climate Change 2007: Synthesis Report", *Intergovernmental Panel on Climate Change*.
- Jackson, M O (2010)**, "Social and economic networks", *Princeton*.
- Jorion, P (2006)**, "Value at Risk: The New Benchmark for Managing Financial Risk", *3rd Edition, McGraw-Hill*.
- Keynes, J M (1936)**, "General theory of employment, interest and money", *Macmillan Cambridge University Press*.
- Kirman, A (2011)**, "Complex Economics: Individual and collective rationality", New York: Routledge.

**Kiyotaki, N (2011)**, "A perspective on modern business cycle theory", Federal Reserve Bank of Richmond Economic Quarterly, 97(3), pp. 195-208.

**Knight, F H (1921)**, "Risk, Uncertainty and Profit", *Houghton Mifflin Company, Boston*.

**Lexis, W (1877)**, *Theorie der Massenerscheinungen in der menschlichen Gesellschaft*.

**Mandelbrot, B (1963)**, "The variation of certain speculative prices", *Journal of Business*, XXXVI, pp. 394-419.

**Markowitz, H M (1952)**, "Portfolio Selection." *Journal of Finance*, 7 (1), pp. 77-91.

**Mirowski, P (1989)**, "The Probabilistic Counter-Revolution, or How Stochastic Concepts came to Neoclassical Economic Theory", *Oxford Economic Papers*, 41(1), pp. 217-235.

**Newman, M E J (2005)**, "Power laws, Pareto distributions and Zipf's law", *Contemporary Physics*, 46 (5), pp. 323-51.

**Noss, J (2010)**, "Extracting information from structured credit markets", *Bank of England Working Paper Series No. 407*.

**Peirce, C S (1873)**, "Theory of Errors of Observations", *Report of the Superintendent US Coast Survey*, Washington, Government Printing Office, Appendix no. 21, pp. 200-224.

**Schularick, M and Taylor, A M (2009)**, "Credit Booms Gone Bust: Monetary Policy, Leverage Cycles and Financial Crises, 1970-2008", *NBER Working Paper Series 15512*.

**Slutsky, E (1927)**, "The Summation of Random Causes as the Source of Cyclical Processes", *Econometrica* (5), 105-146.

**Stigler, S (1999)**, 'Statistics on the Table: the History of Statistical Concepts and Methods', Harvard university press.

**Taleb, N N (2001)**, "Fooled by Randomness: The Hidden Role of Chance in Life and in the Markets, Random House & Penguin.

**Taleb, N N and Blyth, M (2011)**, "The Black Swan of Cairo: How suppressing Volatility Makes the World Less Predictable and More Dangerous", *Foreign Affairs* 90 (3), 33-39.

**Tarashev, N and Zhu, H (2008)**, "Specification and Calibration Errors in Measures of Portfolio Credit Risk: The Case of the ASRF Model," *International Journal of Central Banking* (4), 129-174.

**Turner, S, Farmer, D. J., Geanakoplos, J. (2010)**, "Leverage causes fat tails and clustered volatility", Cowles Foundation Discussion Paper No 1745R.

**Wilson, E B and Hilferty, M M (1929)**, "Note on C.S. Peirce's Experimental Discussion of the Law of Errors", *Proceedings of the National Academy of Sciences of the United States of America*, 15(2), pp. 120-125.

**Vasicek, O (1991)**, "Limiting loan loss probability distribution", KMV Working Paper.

**Vasicek, O (2002)**, "Loan portfolio value", *RISK*, December 2002, pp. 160-162.

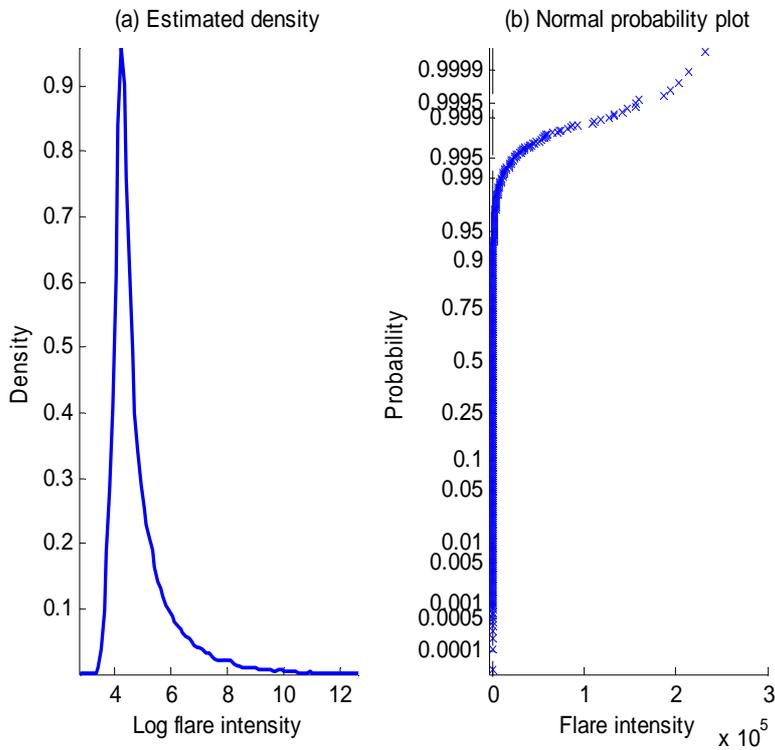
Data annex

The data referred to in the text are drawn from the following sources:

Title	Description	Source	Available at (if applicable)
Rainfall	Monthly total rainfall (mm) in Oxford, UK, January 1853-May 2011	UK Met Office	<a href="http://www.metoffice.gov.uk/climate/uk/stationdata/oxforddata.txt">www.metoffice.gov.uk/climate/uk/stationdata/oxforddata.txt</a>
Rice	Annual percentage change in rice spot price, Bangkok (US \$/Metric Ton), 1011-2011	Global Financial Data	n/a
Equities	Monthly percentage change in stock prices US (S&P500 composite 1791-2012), UK (FTSE All-share 1693-2012), Japan (Nikkei 225 Average, 1914-2012), France (CAC All-Tradable, 1801-2012), Spain (Madrid SE General Index, 1832-2012), Germany (CDAX Composite Index, 1840-2012), Australia (ASX All-Ordinaries, 1875-2012).		
Flares	Peak gamma ray intensity of solar flares between 1980-9.	Newman (2005)	<a href="http://tuvalu.santafe.edu/~aaronc/powerlaws/data.htm">http://tuvalu.santafe.edu/~aaronc/powerlaws/data.htm</a>
Quakes	The intensities of earthquakes occurring in California between 1910 and 1992, measured as the maximum amplitude of motion during the quake. (Magnitudes on the Gutenberg-Richter scale.)		
Cities	The human populations of US cities in the 2000		
Citations	The number of citations received between publication and June 1997 by scientific papers published in 1981 and listed in the Science Citation Index		
Words	The frequency of occurrence of unique words in the novel Moby Dick by Herman Melville		
Loans	Annual real bank loan growth, 1880-2008, UK, US, Australia, Canada, Germany, Denmark, Spain, France, Italy, Netherlands, Norway, Sweden	Schularick and Taylor (2009).	n/a
GDP	Annual real GDP growth, 1880-2008, UK, US, Australia, Canada, Germany, Denmark, Spain, France, Italy, Netherlands, Norway, Sweden		

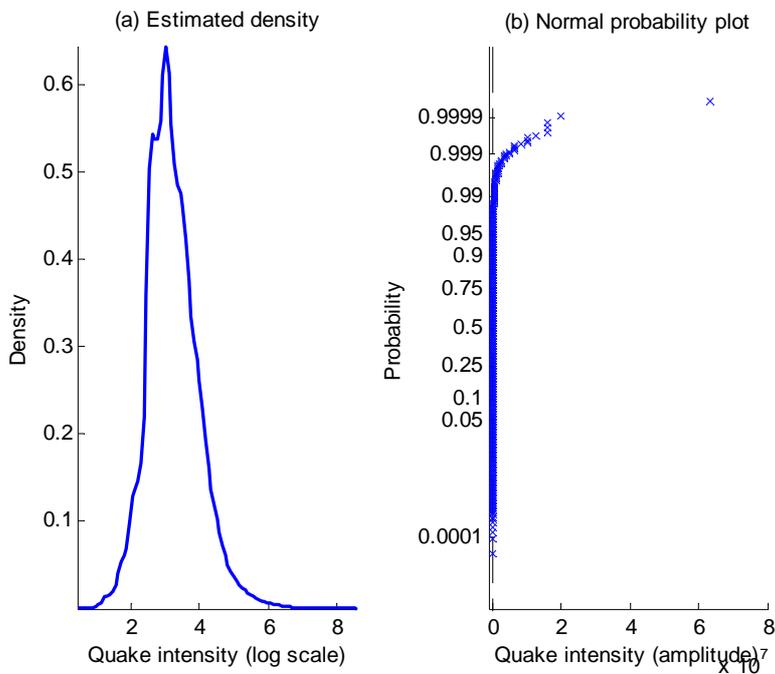
Charts and tables

**Chart 1: Peak gamma-ray intensity of solar flares between 1980 and 1989**



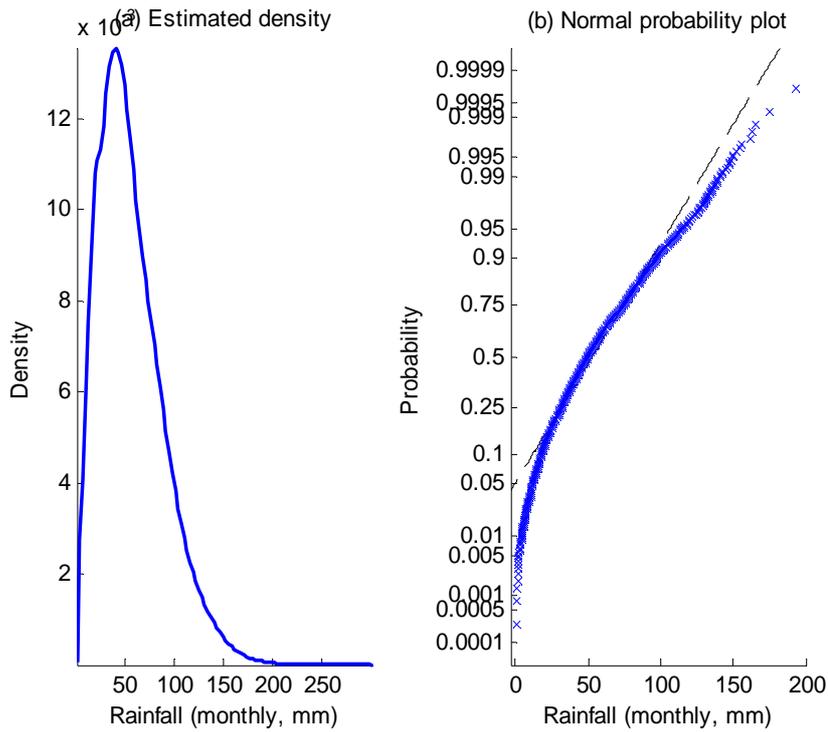
Source: Newman (2005).

**Chart 2: Intensity of earthquakes occurring in California between 1910 and 1992<sup>(a)</sup>**



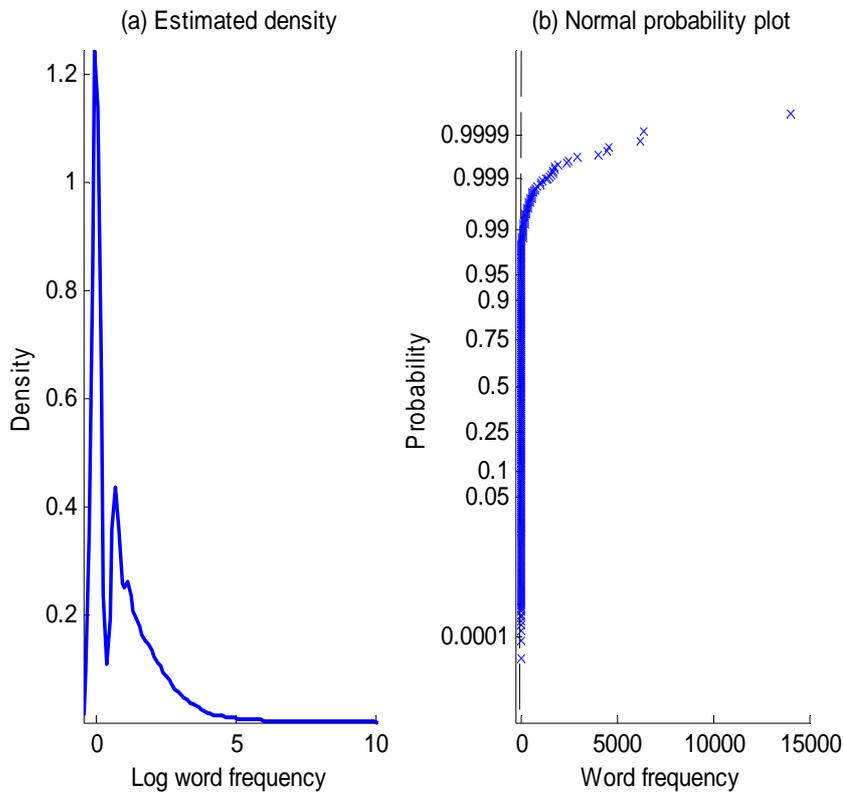
Source: Newman (2005).  
(a) Logarithmic Richter scale.

**Chart 3: Monthly rainfall in Oxford, UK, 1853-2011**



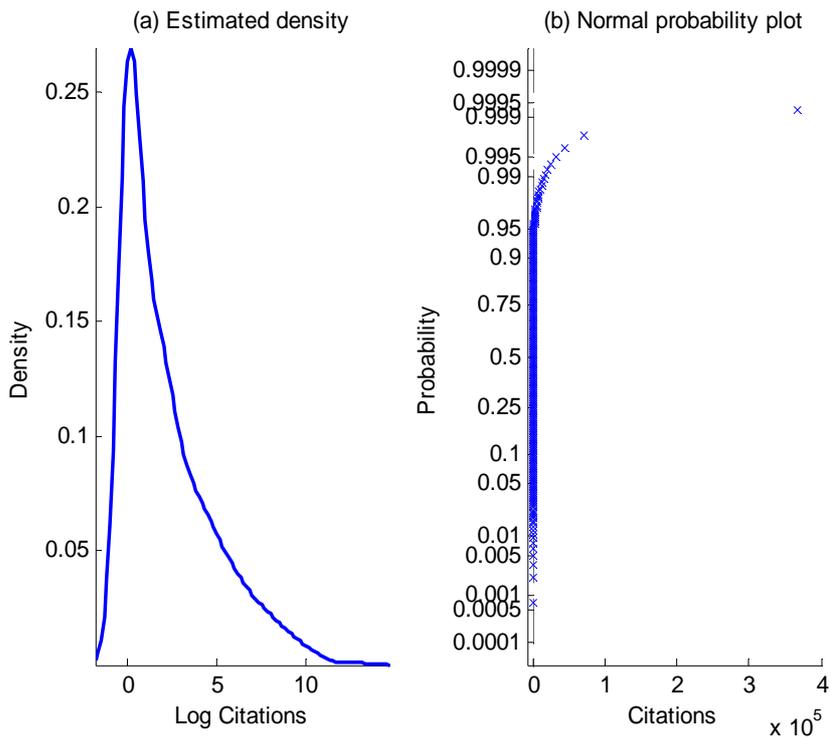
Source: UK Meteorological Office.

**Chart 4: Frequency of occurrence of words in *Moby Dick***



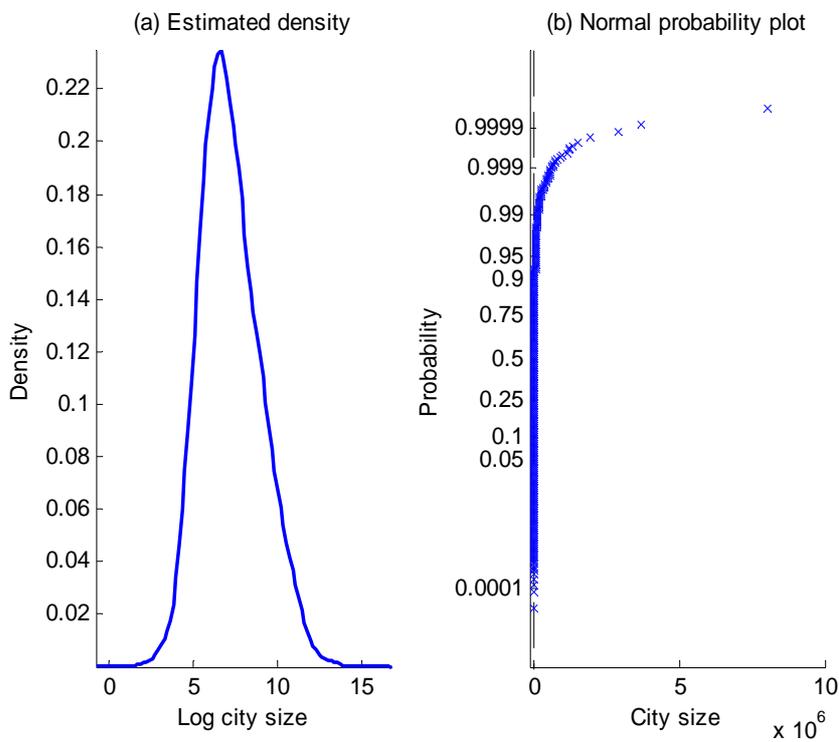
Source: Newman (2005).

**Chart 5: Citations of scientific papers published in 1981, from time of publication to June 1997**



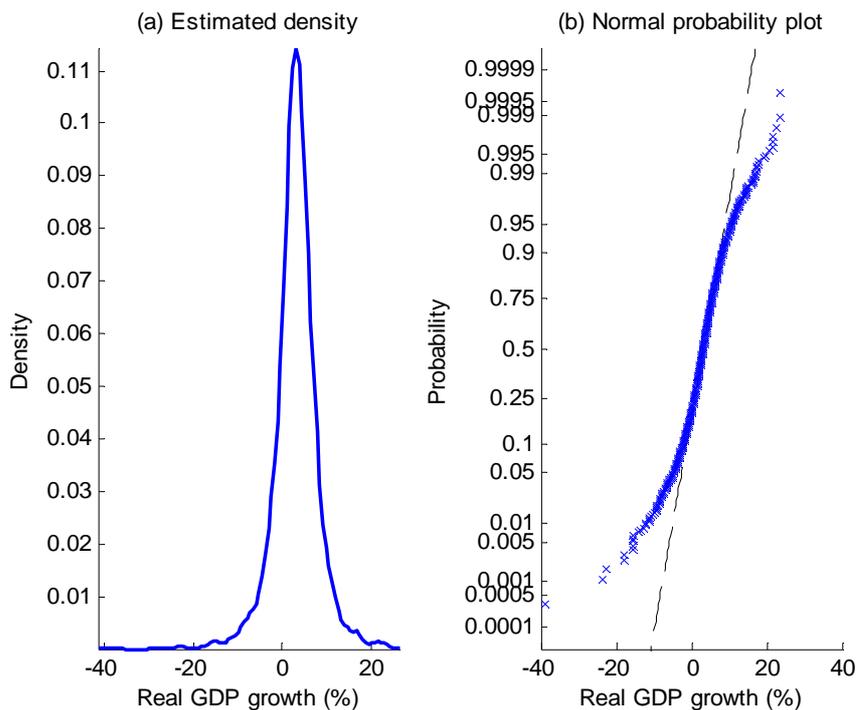
Source: Newman (2005).

**Chart 6: Human populations of US cities, 2000 US Census**



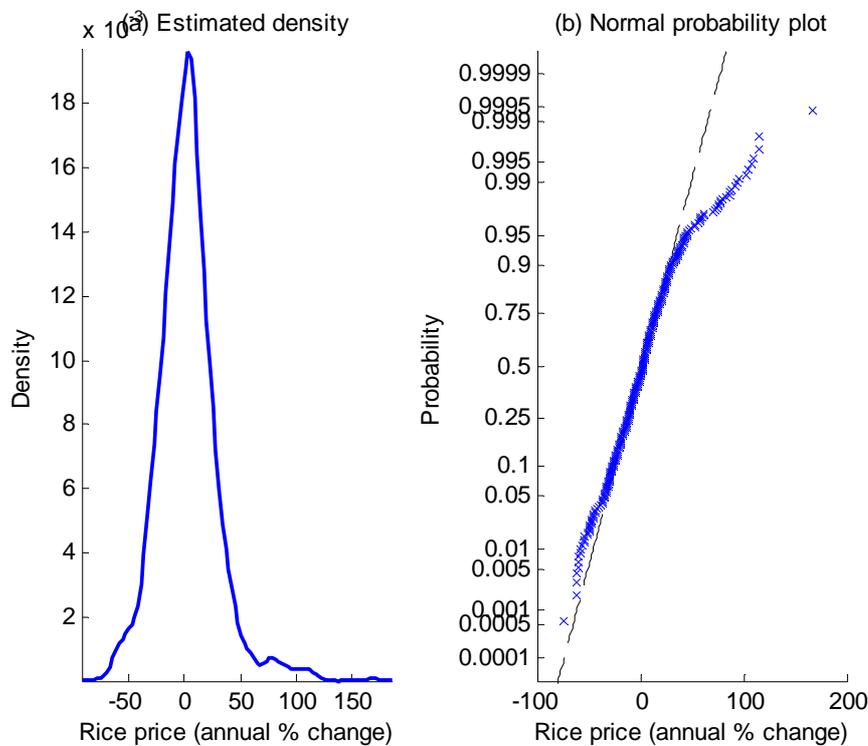
Source: Newman (2005).

**Chart 7: Real GDP growth, 1880-2008**



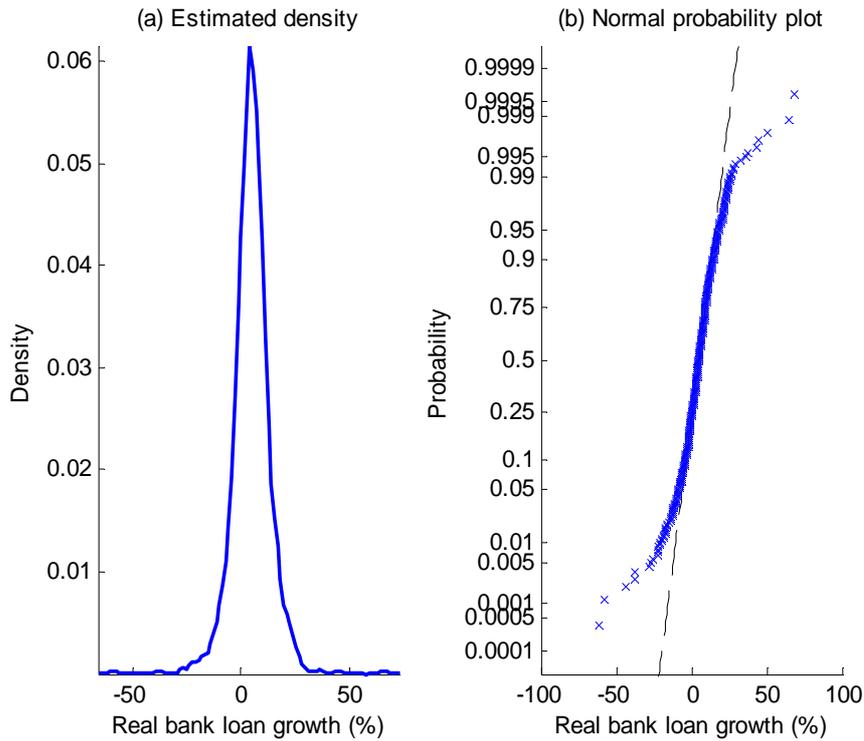
Source: Schularick and Taylor (2009).

**Chart 8: Bangkok rice spot prices, 1011-2011, annual percentage changes**



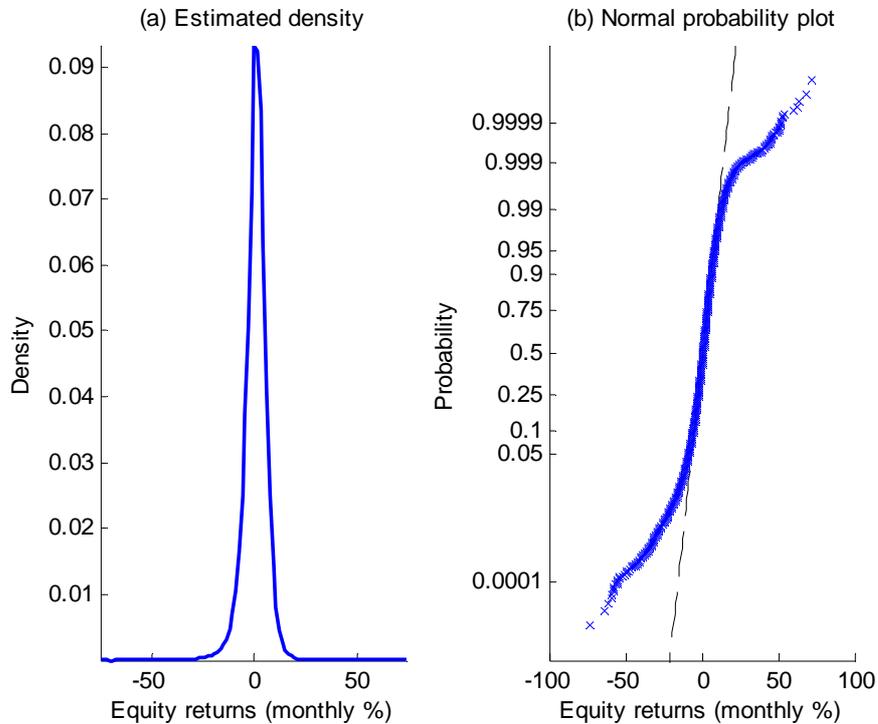
Source: Global Financial Data.

**Chart 9: Bank loans, 1880-2008**



Source: Schularick and Taylor (2009).

**Chart 10: Monthly equity returns, 1693-2012**



Source: Newman (2005).

**Table 1: Kurtosis Statistics for Sample Systems**

	<i>Flares</i>	<i>Quakes</i>	<i>Rainfall</i>	<i>Words</i>	<i>Citations</i>	<i>Cities</i>	<i>Real GDP</i>	<i>Rice</i>	<i>Credit</i>	<i>Equity prices</i>
Kurtosis	585.2	8325.8	3.3	4797.0	606.6	6143.9	9.3	6.8	12.2	12.2

**Table 2: Estimated densities**

	Estimated density	
	<i>Probability of 4 sigma catastrophe<sup>(a)</sup></i>	<i>Ratio to normal<sup>(b)</sup></i>
Probability under normal	0.003%	
<i>Natural</i>		
Flares	0.42%	133
Quakes	0.16%	51
Rainfall	0.07%	21
<i>Social</i>		
Cities	0.28%	88
Citations	0.28%	88
Words	0.22%	69
<i>Economic</i>		
Rice	0.50%	158
Equities	0.43%	136
Loans	0.37%	117
GDP	0.34%	107

(a) A "catastrophe" is defined as an event four sigmas above the mean for the natural and social systems and for rice, and four sigmas below the mean for all other economic systems.

(b) The "ratio to normal" is calculated as the ratio between the estimated density of the four-sigma event for the observed empirical data and the density implied by the normal distribution.