Extreme risks

White paper

2013

TOWERS WATSON

Extreme risks, the irreversibility of time and the retirement anomaly

In short

This paper argues that the world is a complex adaptive system. A consequence of which is that the probabilities of extreme events are higher than typically anticipated. In addition the paper notes that time is irreversible, which means attaching greater weight to the consequences of outcomes and less weight to their likelihood. As a consequence, extreme risks matter and they deserve more attention than they have been given thus far. The paper suggests that retirement for the masses is at serious risk, at least in terms of current expectations regarding length, quality of life, and degree of financial freedom. Alternatively, perhaps retirement as currently configured was never affordable, a reality obscured by demographic and debt trends over the past half century.

Introduction

In the context of the entirety of human history, retirement for the masses is the briefest of anomalies. While the rich have always enjoyed their leisure, the rest of us have worked until we were physically or mentally incapable of carrying on. Assuming we had had enough children so that some of them survived until adulthood we would look to them to support us in our final years. Perhaps our work changed from paid employment to non-paid domestic chores – whether caring for grandchildren, or cooking the family meal, or telling stories and transmitting the family's culture. In any event the idea of self-sufficient leisure time, or 'retirement', would have been an anomaly for approximately 99.9% of human history. It is therefore a very modern invention. The question then becomes, is it a passing anomaly or is it sustainable? To explore this question, we ask whether modern retirement systems are likely to be resilient to extreme risks – low probability, but very high impact shocks. To do so, we evaluate how certain we are in thinking the probability of such shocks is low and how to think about probabilistic modelling in this context.

In what follows, we begin by considering a ‘straw man’ world-view of stability, linear relationships, predictability and equilibrium. Next we introduce an alternative perspective, one of complex adaptive systems with jumps, non-linearities and punctuated equilibria; this perspective offers less predictability and higher likelihood of extreme events. In this second case, it is less possible to diversify risk across time, hence one must weight more heavily the consequence of outcomes. Accordingly, it becomes essential to evaluate extreme risks – very negative and damaging events – so as to identify potential threats to retirement as we know it. We conclude that extreme risks matter and deserve far more attention than given thus far. Partly as a consequence, but also due to Keynes’s ‘paradox of thrift’, retirement for the masses is at risk.

The traditional world-view

Early developers of financial economics relied on economic theory, which itself drew on the mathematics and statistics of physics (Lo and Mueller, 2010). Unfortunately economics borrowed from Newtonian physics with its cause-and-effect linear relationships, and did not go back to borrow from the less-linear world of quantum physics. As a consequence economics is currently stuck in a position of actually being two disciplines; the micro-economics of households and companies maximising and minimising, and the macro- economics of Gross Domestic Product (GDP), monetary and fiscal policies and international trade. Clearly households and companies, with some help from government, create GDP and yet we are only able to study the components or the aggregate, not both simultaneously. This is problem number one. Number two is the assumptions used. For example the economy is assumed to be in, or moving towards, a long-term equilibrium which means we may assume stable parameters through time. We know that the real economy is dynamic, and so short-term cycles are permitted – typically generated by external shocks to the system. These assumptions are a problem because they make the world a very safe place. If we always have an equilibrium, or 'magnetic north', to return to then we can characterise the world with a normal (Gaussian) distribution implying a tight range of expected outcomes and only tiny probabilities of extreme outcomes.

Furthermore, finance has borrowed economics' linear thinking. In economic textbooks or papers we frequently read the phrase '*ceteris paribus*', or 'all other things being equal (constant)', when clearly in a dynamic system they are not and cannot be. In finance this forces us to consider entities (whether companies, portfolios, derivatives, or pricing models) in isolation, and therefore we have no insight on how different market participants, or products, influence each other. Combine normal distributions and linear thinking and it becomes obvious that we know next to nothing about tail events and have no theory of systemic risk. That is, we know the least about the events most able to compromise the sustainability of the retirement system.

The world as a complex adaptive system

To escape the limitations of the traditional world-view we can turn instead to complexity science, an emerging discipline originating at the Santa Fe Institute.[1] Complexity science is the study of complex adaptive systems. Rather than impose on the system our conception of how it works, we can instead study the components of the system and observe how the system-level behaviour emerges from the interactions between those components. All complex systems share some common properties (Mitchell, 2009): first, complex collective behaviour generated by simple individual components, operating with simple rules and no central control; second, there is signalling and information processing (the components produce and use information from internal and external sources, implying that some form of network structure forms the ‘skeleton’ of the system); and third, complex systems are able to adapt and change their behaviour (implying an absence of stable equilibrium through time).

Viewing financial markets as complex adaptive systems rather than as linear systems tending toward equilibrium has important implications (see **Figure 01**). First, understanding each component’s behavior may not provide insight into the system as a whole. Thus, even if all investors were fully rational optimizers, this still may not imply rational, well-ordered and efficient markets. Second, one must grapple with the ‘interconnectedness of all things’. For example, the collapse of Lehman Brothers in September 2008 shows that counterparty risk is subject to non-linear contagion. Third, sudden and violent regime change is possible. Rather than returning to equilibrium after a nasty shock, it may be that turbulence endures for some time. Fourth, and most material to the consideration of extreme risks, is that the tails of the ‘complexity distribution’ are considerably fatter than those of the normal distribution. Complex adaptive systems tend to have scale-free distributions (power laws), so extreme events will be much more likely. Fifth, modelling becomes much more complex. That is, shocks can be modelled with fatter tails and higher probabilities, but complexity science has not yet developed sufficiently to provide readily-usable quantitative models for finance and economics.

Figure 01. Some characteristics of complex adaptive systems

Characteristic[a]	**Financial market**[b]
Underlying simplicity	An individual buys, or sells, a security.
Many components	Many investor types, many intermediaries, exchanges, payment systems.
Many individual actors	Each major market has thousands of institutional investors, and multiple thousands of individual investors. The investors span markets resulting in inter-connectedness of apparently separate systems.
Multiple spatial and temporal scales	Stock price charts over different lengths of time cannot be distinguished, meaning that stock price movement is scale-free with respect to time. It is therefore better proxied by a power law distribution than a normal distribution. This is consistent with stock price movements being self-similar (fractals) (Mandelbrot 2004).
Unintended consequences	The repeal of the Glass-Steagall Act, which separated retail and commercial banking, was justified on the grounds of increased financial sophistication and new risk management technology. However, in retrospect it allowed deposit-taking banks to leverage their balance sheets thereby contributing to the severity of the global financial crisis.
Emergent phenomena	Bubbles and busts are unintended, and are not controlled by any single actor, or group. Consequently they can be considered as a behaviour of financial markets that emerges from multiple individual interactions.
Historically contingent /path dependent	Bond market returns over the last 30 years were heavily dependent on yields falling from around 14% to around 2%. If we agree that yields will not fall another 12% over the next 30-years (to -10%) then it follows that bond returns cannot be as good in the next 30 year period. Future bond returns are historically contingent on what happened in the prior period. If

Characteristic[a]	Financial market[b]
	returns were independent and identically distributed, the returns over the next 30 years could be the same as the previous 30. In equities, the concept of path dependency can be illustrated by Soros's concept of 'reflexivity' where "the mispricing of financial instruments can affect the fundamentals that market prices are supposed to reflect" (Soros 2012).
Multiple phases (regimes)	Academics have identified different market regimes in monthly returns data (Kritzman 2010) and over spans of a decade or more (Brock 2003). The difference in these timescales is interesting and refers back to the multiple time scales characteristic above.
Non-linear	Most, if not all, bubbles and busts would qualify as non-linear market events. However, there is an interesting example in the quantitative equity crisis of August 2007. David Viniar, then Chief Financial Officer of Golman Sachs, explained that 'we were seeing things that were 25-standard deviation moves, several days in a row'[c]. This quote inspired academic papers exploring the implications. In short, a 25-sigma event would be so incredibly rare that most other things would be more likely. My argument is that the confluence of quantitative equity trades caused prices to move in a non-linear way.
Robust/resilient	The global financial crisis notwithstanding, it is remarkable given the volume and size of daily transactions, as well as the ongoing attempts by unethical players to pervert markets for their own ends, that more doesn't go wrong more often.
Adaptive/ evolving	If we compare financial markets today with any previous point in time we will see remarkable changes and yet seldom have these changes appeared material at the time. The underlying simplicity remains – the buying and selling of securities – but now the trades are in decimals, occur in fractions of a second, and may not happen on an exchange. In addition the number and variety of securities available to transact has ballooned, particularly derivatives.

Characteristic[a]	Financial market[b]
Non-equilibrium	There is no proof that can be offered to demonstrate conclusively that financial markets do not have an underlying equilibrium. Ultimately it comes down to an individual's belief about how the market operates. However, the clear exhibition of the other characteristics suggests, to the author at least, that financial markets do not exist in equilibrium. Instead they exhibit 'punctuated equilibria' where they appear to be in equilibrium while in fact cumulative change is occurring 'beneath the surface' which at a certain point leads to a non-linear jump. Hyman Minsky's financial instability hypothesis is an example of a punctuated equilibrium model.
Complicated vs complex	An aeroplane is complicated as it requires many different parts and systems to come together in a precise way in order to be able to fly safely. However it is not complex. Once those parts and systems have been assembled according to the design the plane will fly predictably. In contrast a complex system is non-repeatable and unique, in that slightly different interactions will lead the system down a different path. While not possible to rewind time and prove by experimentation, financial markets dues to their reflexive nature appear more likely to be complex than complicated.

[a] This column lists the characteristics of complex adaptive systems as given in a presentation by Geoffrey West, Distinguished Professor and Past President, Santa Fe Institute to the Foundational Questions Institute,

[b] This column indicates how financial markets exhibit the characteristics of complex adaptive systems

[c] Financial Times August 13, 2007

Source: http://fqxi.org/data/documents/conferences/2011-talks/west.pdf and author.

To summarise the argument of the paper so far, when it comes to assessing risks, particularly low-probability, high-impact events, our limited understanding of the world can have a material impact. In fact, the uncertainty and our proneness to error can dominate when the extreme events involve poorly understood natural phenomena, complex social dynamics such as financial markets, or new technology (Bostrom 2013). For example, suppose that our body of knowledge indicates that some catastrophic event X has an extremely low probability Pr(X) of occurring. The probability that our body of knowledge is flawed may be significant, resulting in a large margin of error. If this seems a strange concept at first, consider that our body of knowledge once thought the solar system was geocentric. In fact the whole history of scientific progress is one of correcting flaws in the previous body of knowledge. Our argument so far is that complexity thinking represents an improvement in our body of knowledge. The downside is that extreme events are likely to be more frequent than we previously thought.

While the discovery of higher probabilities is material in its own right, there is a further reason why we should pay greater attention to extreme risks, namely that we only get to live once.

The irreversibility of time

As strange as it may appear much of finance and economics implicitly assumes we have infinite lives all running in parallel. To illustrate the point, consider the following gamble. You will roll a fair dice and if you roll any number from one to five we will pay you 50% of your current wealth, including the present value of your future earnings. This is a thought experiment so we will gloss over our ability to pay – assume our credit is pristine. Imagine how much better your life would be if you were one-and-a-half times richer in the time it took to roll a dice. The downside, paltry in comparison, is that if you roll a six you will pay us your entire wealth – house, pension pot, all future earnings, the lot. Will you take the gamble?

The way we have been trained to analyse the gamble means that we will consider all the possible future outcomes and then weight them in accordance with their probability. In effect we freeze time and take multiple copies of the world and then run the six versions forward as 'parallel universes'. In one of those worlds a one is rolled and we pocket a 50% gain in wealth. In the second a two is rolled with the same result. In the sixth world a six is rolled and we lose all our wealth. Having exhausted all the possibilities we travel back in time to the present and do our sums. The expected return of the gamble is the ensemble average – the average of all the possible independent outcomes. In this case the expected return is 25% and so we would be 'crazy' not to take it.

So would you take the gamble? The answer is typically 'no'. Instinctively, something does not feel right. Either you do not trust our credit, or the ensemble average (expected return or expected value) is misleading in some way. So let us consider the time average instead. Instead of rolling the dice once in each of six parallel universes, we will stay in our familiar universe and roll the same dice six times in succession. We compute the time average by taking each of the six possible independent outcomes and making them occur one after the other in our single, real, universe. We now compound our returns over the six periods and take the sixth-root to calculate our per-period expected (time average) return. It does not matter what order we roll each of the numbers one to six, we will lose all our wealth and so the time average is negative, and in a big way (-100%). So the ensemble average *is* misleading. The 25% expected return unhelpfully disguises the meaningful (16.7%) likelihood that we lose everything.

The point of this thought experiment is to introduce the notion that we cannot go backwards in time, as once we have lost everything we cannot go back and try again. The more subtle point is that the traditional calculation we use, the expected return (ensemble average), effectively underweights the significance of extreme risks (for more detail on the impact of the irreversibility of time please see Towers Watson (2012a)). We can illustrate this with a more realistic thought experiment. Consider a world, in which we are investing our portfolio of financial assets, where there are two types of outcomes; good outcomes, which produce a return of 5% and occur almost all the time, and extreme outcomes which only occur once in the distribution but cause severe or total loss. If the ensemble average truly does understate the significance of the single extreme event then it will consistently overestimate the likely return our portfolio will achieve relative to the return the time average suggests we will achieve.

Figure 02 shows that this is indeed the case for a number of 'runs' of our thought experiment world. The pairs of columns represent runs for worlds with different probabilities for the single extreme outcome starting with 1-in-1,000 and moving right to 1-in-100. The rows also show different worlds, where the severity of portfolio loss increases from 99% in the first row to 100% in the bottom row. Note that for all 20 runs (combinations of probability and severity) the ensemble average return is always higher than the time average return, and in some cases significantly higher. In addition, please note that once the probability of the extreme event gets up to 1-in-100, or higher, then 99 good outcomes of 5% are wiped out by a single extreme event. Finally, note the difference in the time average return between extreme losses of 99.999% and 100%. Like our artificial dice-throwing experiment above, once you have lost all your wealth the game is over and your return is -100%, whether that occurs in the first period or the last period. Losing 99.999% of your portfolio would clearly be painful, but the little that is left can then start to grow again.

Essentially this is highlighting the difference between an existential risk and a risk where 'life' continues into the next period, albeit in very poor shape.

Figure 02. The expected percentile return from a distribution comprising good outcomes with a 5% return and a single extreme outcome

	Probability of extreme event occurring:							
	1-in-1000		1-in-500		1-in-250		1-in-100	
Loss given extreme event	Ensemble average	Time average	Ensemble average	Time average	Ensemble average	Time average	Ensemble average	Time average
-99.000%	4.90	4.51	4.79	4.03	4.58	3.06	3.96	0.23
-99.900%	4.90	4.27	4.79	3.55	4.58	2.12	3.95	-2.06
-99.990%	4.90	4.03	4.79	3.07	4.58	1.18	3.95	-4.29
-99.999%	4.90	3.79	4.79	2.60	4.58	0.25	3.95	-6.46
-100.000%	4.90	-100	4.79	-100	4.58	-100	3.95	-100

Notes:

The ensemble average is calculated using the formula

$$\mathrm{E(r)} = \sum_{i=1}^{N} \frac{r(i)}{N}$$

The time average is calculated using the formula

$$T(r) = \prod_{i=1}^{N} [1 + r(i)]^{1/N} - 1$$

Source: Author's own calculations.

Now it is possible to object that a loss on a portfolio of 99% or more is too extreme to realistically contemplate. The point of a portfolio, after all, is to diversify against such extreme losses. Clearly this is a valid objection; however this is a thought experiment so the value is in what it teaches rather than the realism. That said, we believe that some of the extreme risks that we discuss below could indeed cause portfolios of financial assets to become worthless. Besides, there are several historical examples of entire stock market losses. Returning to the learning point though, rather than consider a literal 99% loss (or greater) instead consider what effect a large portfolio loss could have on a retirement fund. Here you can mentally adjust 'large' as you see fit, but perhaps start with a 50% loss and adjust higher and lower. If the retirement fund is a defined benefit arrangement and the sponsoring employer is now small relative to the fund, or has ceased trading, then we would argue that a large, and feasible, portfolio loss can represent an existential event in that context. By 'existential' in this instance we mean that the mission of the retirement fund will have failed at that point. The assets will run out before the liabilities are paid and absent an insurance arrangement, some of the beneficiaries will receive nothing. So for them at least this would equate to a total portfolio loss.

If instead the retirement fund was a defined contribution arrangement then it is less likely that the large portfolio loss would qualify as existential. This is due to the fact that there is no contractual benefit to be broken. Instead the members 'get what they get' and the adjustment mechanism is up to the individual member, perhaps by accepting a lower standard of living in retirement than hoped for. Even here, however, not all members are equal. A 50% loss for a 29-year old is fundamentally different to a 50% loss for a 59-year old, and the older member may be tempted to consider that they had suffered a loss bordering on the existential. The practical takeaway is that avoiding, or reducing the probability of 100% (existential) losses is incredibly valuable and should become a top priority.

We have established that using a different lens through which to look at the world can significantly increase the qualitative probability of an extreme event occurring.[2] And we have seen that embracing the reality of irreversible time meaningfully increases the significance of extreme risks. Having laid that ground work we can now move on to consider the events that we might label as extreme.

Extreme risks identification

We have already noted that extreme risks are very unlikely but high impact events. There is a potentially very long list of risks to consider and so we must develop a framework, or set of filters, to help us to identify which risks we should assess, and which we can safely ignore. For ease of exposition we have chosen to consider only first-order impacts, and to only consider risks in isolation. In a specific risk-management situation both of these constraints should be removed. For example, a German reinsurance company undertook an analysis of the risks faced by each of its lines of business and developed a new software tool to allow for subsequent-order effects. The results showed that the obvious risks were indeed the biggest risks that the company faced, but in the top five was a surprise – a European heat wave. The heat wave did not feature as a first-order risk but was prevalent across multiple lines of business. In this example, German law prohibits power stations from drawing water from, and discharging waste water into, rivers once the river temperature crosses a threshold. This means power stations must shut down, impacting electricity users and triggering insurance claims. Such analysis is important in an individual context but beyond the scope of this paper. However an attempt is made to move towards second-order considerations in the form of an 'association matrix', a qualitative assessment of whether there is likely to be a causal link between the individual risks. This is expanded upon below.

The second constraint we have adopted is to consider the risks in isolation, and much the same comments apply. Context-specific risk assessment should consider whether a chain of events could occur; whether the conditional probability of a subsequent event is substantially higher than the unconditional probability; and whether the combined events have a different scale of impact. As an example, we would expect German luxury car manufacturers to be unaffected by an earthquake in the Pacific. However if that earthquake were to cause a tsunami that hit the coast of Japan then maybe we would factor in a small drop in sales of luxury cars. But if that tsunami causes a nuclear accident which closes down Japanese manufacturers then the production of luxury cars in Germany can be halted for want of a single component. And halted for as long as it takes to rejig the global supply chain. Clearly the clustering of risks such as this is the basis of scenario analysis which should be done but is outside the scope of this paper.

An additional filter that we will apply is that of context. In this paper we are concerned about the sustainability of the retirement system and so we focus primarily on risks that will diminish the asset side of the balance sheet or increase the liability side. As a consequence something like extreme longevity becomes an extreme risk whereas in most other contexts this would be seen as a positive. In a similar way our consideration will exclude legal, process and operational risks.

Having acknowledged these constraints we now outline the overall framework.

Framework. The framework we have used consists of six categories; finance, economics, politics, environment, society and technology. The six categories are not independent. In fact the most interesting insights are often found at points of confluence. Picture the categories as small overlapping circles at the centre of a page. It is then possible by free association to sketch out a 'mind map' of drivers, issues and risks that interlink and can often be influenced by multiple categories. The downside of mind maps is that they are most useful to the originator as they summarise the thought process while leaving out the detail. Furthermore, experience teaches that most people are more comfortable reviewing a linear list and for both these reasons this is the approach used in this paper.

The extreme risks covered in this paper are shown in **Figure 03** together with a brief description (for our earlier treatments of extreme risks please see Towers Watson (2010, 2011)). To be clear, some of these risks are likely to seem very extreme (perhaps amusingly, or ridiculously, so). But there is a serious point here, namely how we should try to deal with black-swan-type events, particularly as human nature appears to think that if an event has not occurred it cannot occur. Please note that within each category the risks appear in alphabetical order and do not represent any form of ranking at this stage. For a fuller account of each risk please refer to the appendix. The following parts of this section discuss each of the four categories in turn. To restate the context of this paper the risks are intentionally extreme. In effect we are asking 'is it plausible?' or, perhaps more accurately, we are asking 'does our current state of knowledge suggest it is completely implausible?'. Having made it on to the list we will subsequently ask 'assuming this event occurred, what would be the consequences?'.

Figure 03. The extreme risks by category

Financial	
F1 Banking crisis	Central banks unable/unwilling to supply liquidity in the next crisis, causing banking and real economic activity to stop.
F2 Insurance crisis	Insolvency of the insurance industry through inadequate capitalisation of risky asset positions or through 'basis' risk.
F3 Sovereign default	Non-payment by a major sovereign borrower causes market panic and adversely disrupts the global economy.
Economic	
E1 Abandonment of fiat money	A collapse in confidence in the purchasing power of paper currency and the consequent return to a gold standard.
E2 Break-down of capitalism	Distrust in the private capital/property system, causing a collapse in economic activity and asset prices.
E3 Currency crisis	A significant devaluation of a major currency that becomes self-fulfilling, with loss of purchasing power.

E4 Deflation	Goods and services prices fall for a long period, transferring wealth from borrowers to savers; often associated with depression.
E5 Depression	A deep and protracted trough in economic output, massive increase in unemployment, restriction of credit, shrinking investment.
E6 Hyperinflation	Prices increase rapidly, wiping out savings, provoking extreme consumption and hoarding of real assets.
E7 Stagnation	A prolonged period of little or no economic growth, usually accompanied by high unemployment.
<u>Political</u>	
P1 Anarchy	An extreme form of social disorder in a major economy, resulting in a loss of power by government and possible imposition of martial law.
P2 Global trade collapse	Extreme rise in protectionism causing global trade and investment to collapse, requiring a rise in self-sufficiency.
P3 Political extremism	The rise to power in a major economy of a brutal and oppressive government, typically causing a large number of civilian deaths and becoming a major threat to global peace (for example, totalitarianism, whether fascism or Stalinism).
P4 Terrorism	A major (ideologically-driven) terrorist attack, targeted at a region of global economic and/or political importance.
P5 World War III	A military war involving many of the world's most powerful and populous countries causing multiple-millions of deaths (could involve a nuclear holocaust).
<u>Environmental</u>	
e1 Alien invasion	Invasion of non-peace-seeking extra-terrestrial beings that either removes the planet’s resources or enslaves or exterminates humanity.
e2 Biodiversity collapse	The destruction of the world’s ecosystem and therefore the loss to humans of ecosystem services: provision (food and clean water), regulation (climate and disease), support (nutrient cycles and crop pollination) and culture (spiritual and recreational benefits).
e3 Cosmic threats	Existential risks arising beyond earth, such as a major meteorite impact, being pulled out of orbit (or the solar system) by a passing asteroid, or a giant solar flare (which would be compounded if during a reversal of the earth’s magnetic field).

e4 Global temperature change	Habitable areas in the earth are significantly reduced due to excessive heat or cold, associated with significant sea level rises or 'ice age' respectively.
e5 Natural catastrophe	Earthquakes, tsunamis, hurricanes, flooding (including atmospheric river storms) and volcanic eruptions. The extreme risk would either be a confluence of connected extreme natural catastrophes or the eruption of a supervolcano (ejecta >1,000km^3 threatening species extinction).
Social	
S1 Extreme longevity	An unanticipated, significant increase in life expectancy for much, or the majority, of humans. Possibly the result of a major breakthrough in medical or human genome science.
S2 Food/water/energy crisis	Major shortfall in the supply of food, water, or energy causing political strife and widespread death and severe damage to the quality of life for many survivors.
S3 Health progress backfire	A reversal in the trend of improved health. Possibly caused by societal trends such as mental health problems or obesity, the unintended consequences of current healthcare practices (eg antibiotic resistance), or a slowing in the rate of medical advancement below the rate of pathogen evolution. Economic output falls and liabilities increase.
S4 Organised crime	A significant increase in the scale of illegal operation in a major economy or region to the extent that legitimate economic activity becomes non-viable.
S5 Pandemic	An epidemic of highly infectious and fatal disease that spreads through human, animal or plant populations worldwide.
Technological	
T1 Biotech catastrophe	An instance of error or terror causing widespread deaths. The risk arises from the easy synthesis of DNA, as well as other biological manipulations, being increasingly available to small groups of technically competent and even individual users.
T2 Cyber warfare	Politically-motivated computer hacking to conduct sabotage and espionage on a national or global-power scale.

T3 Infrastructure failure	An interruption of a major infrastructure network for a relatively long period due to, either human behaviours, natural disasters, or even cosmic threats.
T4 Nuclear contamination	A major nuclear accident or attack that leads to lethal effects to individuals and large radioactivity release to the environment.
T5 Technological singularity	An extreme risk resulting from technological advancement proceeding beyond the point of human understanding. The creation of a computer more powerful than the human brain, which can then design and build an even more advanced machine creating an environment where human survival is at risk.

Note: The risks are listed alphabetically within the categories.

Source: Author.

Financial. This category of extreme risks concerns an inability to meet liabilities. This can occur as a liquidity event, such as a banking crisis, where an institution has insufficient cash or other liquid assets to meet a current demand for payment, even if the institution has more assets than liabilities. Failure to make that payment can then cascade rapidly through the financial system, with further institutions then unable, or unwilling due to a collapse in trust, to meet their own payment demands. The alternative mechanism for a financial extreme risk is a solvency-driven crisis, such as an insurance crisis or sovereign default, where there are insufficient assets to meet the liabilities irrespective of how liquid the assets are. This does not preclude the possibility that a solvency crisis could lead to a liquidity crisis, or vice versa.

Economic. These risks arise from a shock to growth, a shock to price levels, or a collapse in trust which is essential for the efficient working of any economic system. Growth shocks can take the form of a depression or stagnation. The former has a painful contraction phase but then relatively swift recovery, compared to a stagnation episode where growth is weak for an extended period. Price level shocks can occur in opposite directions; rapid rises in hyperinflation and falling prices in deflation. In both cases the ‘incorrect’ price signals cause serious economic damage and destruction of wealth. A collapse in trust could occur in the current monetary system (abandonment of fiat money), in the value of a major currency (currency crisis), or in the economic system as a whole (break-down of capitalism).

Political. This category of extreme risks comprises those which derive from policy decisions. In two of the cases the link is direct and obvious. Global trade collapse follows policy decisions to favour protectionism over openness and globalisation, and World War III follows an active decision to declare war. For anarchy and political extremism the link is less direct but in both cases poor prior policy decisions are likely to be a necessary – if not sufficient – condition for these risks to foment. Terrorism is included in the political category due to its ideological foundation and as the target chosen for the act of terrorism is likely to have political ramifications. Please note that we are considering extreme manifestations in this paper. Terrorism is a weekly, if not daily, occurrence somewhere around the world and so the extreme risk would be a terrorist act comparable to, or worse than, 9/11.

Environmental. The risks in this category are threats to human safety and well-being arising from a disruption to planet earth’s environment. If we draw the boundary of the system around the earth and its atmosphere then two of these risks – alien invasion and cosmic threats – would be exogenous. Is an alien invasion too extreme to spend any time seriously considering? Quite possibly. After all, both the probability of the event and the consequences are unknowable. However, risk management is about taking action in advance to prepare for possible future consequences. After the event, offering ‘we didn’t think about that’ as an excuse seems to represent a risk management failure. So the value of the exercise is in scanning the horizon with an open mind. We can always apply further filters at a later stage to protect our finite risk management resources. For now all we can conclude is that we do not know enough to rule out the possibility of an alien invasion.

Two of the environmental risks, biodiversity collapse and global temperature change, *could* be caused by humanity, and would thus represent serious own-goals. However there are other possible causes for both of these risks, and these can be reviewed in **Figure 06** below. The final risk in this category is natural catastrophe. As earthquakes, for example, happen every day the extreme version of this risk is either a confluence of extreme natural catastrophes (think magnitude 10 earthquake, combined with 25 metre tsunami, helped along by a category five windstorm) or the eruption of a supervolcano. This is simply the downside of living on a planet that regularly brings to the surface useful and valuable minerals.

Social. The social extreme risks are those threats that could adversely affect the smooth functioning of society. As already noted above, the categories we are discussing are not independent and it should be clear that the social risks link to policy decisions, the environment, and in some cases, to technology. This is obvious in the case of food/water/energy crisis which will have political, environmental and technological drivers as well as offsets.

Three of the risks are health related. Pandemics are a favourite of commentators on extreme risks as in relative terms there is plenty of good data. For our purposes we postulate a new disease agent that hits the ‘disease sweet spot’ of high infectivity and high mortality (these are typically trade-offs). Health progress backfire refers to a reversal in the trend of improved health while, in the other direction, extreme longevity becomes a risk when viewed through the lens of a retirement provider. In most other contexts it would be considered a boon. The final risk in this category is the growth in organised crime to the extent that legitimate economic activity ceases to be viable in the (major) country or region concerned.

Technological. Our final category of extreme risks concerns technology. These risks range from a failure in current technology (nuclear contamination and infrastructure failure), through the possible consequences of emerging technology (cyber warfare and biotech catastrophe), to the unknowable future event of the technological singularity. The latter risk refers to the point in time when humans have designed super-intelligence into machines. What happens beyond that point is unknowable and therefore the subject of speculation. The extreme version has already been foreshadowed in various fictional films where the machines replace their human creators.

Extreme risks assessment

For this exercise to be useful we cannot stop at the identification and simple listing of the risks. We need to assess them in order to determine which are more material and which are less. The preferred approach in risk assessment is to review the historic data and see what we can learn from it. A qualitative overlay can be added if desired, for example if a structural change is believed to have occurred, but the better the data the less this tends to happen. When dealing with extreme risks however, by definition we are dealing with very small data sets at best. As a consequence this paper takes a deliberately qualitative (deductive) approach to assessment. At first glance this may seem peculiar, but we believe the output presented in this paper is significantly better than one individual's guess, and is likely to be better than anything produced by a quantitative approach. It is better because of the process adopted and this is described next.

Process. We followed a rigorous and robust process to develop our qualitative assessment methodology. In stage one, a team of Towers Watson researchers reviewed the research literature and reviewed historical data on past extreme events. The team members then independently generated their scores. In stage two, the independent scores were compared and debated, with a single consolidated scoring approach being the outcome. For stage three, the consolidated scores were sent for peer review by a senior committee, and further refinements were suggested by the committee. Stage four was sign off of the revised scores by the peer review committee. Consequently the results presented in this paper have been subject to a rigorous governance process and are therefore robust despite being qualitative.

Scoring. The scoring of the extreme risks is shown in **Figure 04** and is illustrated graphically in **Figure 05**. Each risk has four scores, for likelihood, uncertainty, intensity of impact, and scope of impact. For the grading of the intensity and scope of impact we have drawn on, and adapted, the qualitative risk categories of Nick Bostrom (Bostrom 2013).

The scores relate to the most extremely negative manifestation of the risk. For example, the consequences of earth being visited by aliens could span a massive range from the beneficial through to the extinction of the human species. Given that we are concerned with extreme risks it is the latter potential outcome that is scored.

Rather than attempt to give each risk a probability (the quantitative route), we assign them one of four likelihood categories representing a likelihood of occurrence of 1-in-10 years, 1-in-20 years, 1-in-100 years, and less likely than 1-in-100 years. For the technically minded this can be thought of as akin to a high-alpha power law for the distribution and implies we believe these events should be considered far more likely than if a normal distribution were used (per the complexity world-view discussed above).

We have split the potential impact of the risk into two separate dimensions, namely the intensity and the scope (or geographical and temporal spread). The intensity is assigned to one of three states that are labelled ‘endurable’, ‘crushing’, and ‘existential’. We do not define these states precisely as we believe the labels are self-explanatory. However, if it helps, consider yourself in the three states. An endurable risk could represent a broken leg; crushing might imply the loss of a limb, or paralysis; and existential could refer to the loss of self-awareness or loss of life. The scope of the impact attempts to convey both spatial and temporal information by use of four categories: local, global, trans-generational, and pan-generational. The first two imply a temporary impact while the latter two imply a lasting impact. We use 'trans-generational' to describe an impact that will affect more than one generation but that would then fade or reverse. Pan-generational is used to describe an impact that would affect all subsequent generations, or all previously potential generations (such as extinction of the human species). We have split the impact into these two dimensions as we believe it gives valuable extra information. However, there is a danger that the dimensions are not completely independent. For example, a food/water/energy crisis could validly be described as either locally crushing or globally endurable. For the majority of other cases we believe the two dimensions are sufficiently independent to provide useful additional information.

The final score assigned to each risk is uncertainty which is assessed as low, medium, or high. In the graphical representation (see **Figure 05**) this is shown as a semi-transparent border around the shape, with higher uncertainty shown by a larger 'fuzzy' border (or 'location'). As indicated by the shapes, the uncertainty is in two dimensions: uncertainty regarding the likelihood and uncertainty regarding the impact. An extra layer of sophistication could have been used at this point by extending the shapes more in either the vertical or horizontal direction, implying greater certainty in one dimension and less in the other. However, if going down this route, why stop there? Why not also move the border off-centre to show that the uncertainty is asymmetric? Given the intrinsic high levels of uncertainty involved when considering extreme risks at all, we concluded that this extra sophistication would represent spurious accuracy, and give a false impression of the level of signal available in this area relative to the noise.

Figure 04 Extreme risk scoring

	Likelihood 1 (1-in-10 years) 2 (1-in-20 years) 3 (1-in-100 years) 4 (1-in-100+ years)	Uncertainty A degree of: (H) High (M) Medium (L) Low	Impact intensity 1 (Endurable) 2 (Crushing) 3 (Existential)	Impact scope 1 (Local) 2 (Global) 3 (Trans-generational) 4 (Pan-generational)
Financial				
F1 Banking crisis	2	L	1	2
F2 Insurance crisis	3	L	1	2
F3 Sovereign default	2	L	2	1
Economic				
E1 Abandonment of fiat money	3	M	1	2
E2 Break-down of capitalism	4	M	2	3
E3 Currency crisis	2	L	1	2
E4 Deflation	2	L	1	1
E5 Depression	2	L	2	2
E6 Hyperinflation	3	M	1	1
E7 Stagnation	1	L	1	1
Political				
P1 Anarchy	3	M	2	1
P2 Global trade collapse	1	M	1	2
P3 Political extremism	3	H	2	1
P4 Terrorism	2	M	1	1
P5 World War III	3	M	2	2
Environmental				
e1 Alien invasion	4	H	3	4
e2 Biodiversity collapse	3	M	2	3

	Likelihood 1 (1-in-10 years) 2 (1-in-20 years) 3 (1-in-100 years) 4 (1-in-100+ years)	Uncertainty A degree of: (H) High (M) Medium (L) Low	Impact intensity 1 (Endurable) 2 (Crushing) 3 (Existential)	Impact scope 1 (Local) 2 (Global) 3 (Trans-generational) 4 (Pan-generational)
e3 Cosmic threats	4	M	3	4
e4 Global temperature change	2	L	2	3
e5 Natural catastrophe	4	M	2	3
Social				
S1 Extreme longevity	3	L	1	2
S2 Food/water/ energy crisis	1	L	2	1
S3 Health progress backfire	2	M	1	3
S4 Organised crime	2	M	1	1
S5 Pandemic	2	H	2	2
Technological				
T1 Biotech catastrophe	3	H	2	2
T2 Cyber warfare	2	H	1	2
T3 Infrastructure failure	2	M	1	1
T4 Nuclear contamination	2	M	2	1
T5 Technological singularity	3	H	3	4

Figure 05. Graphical representation of risks

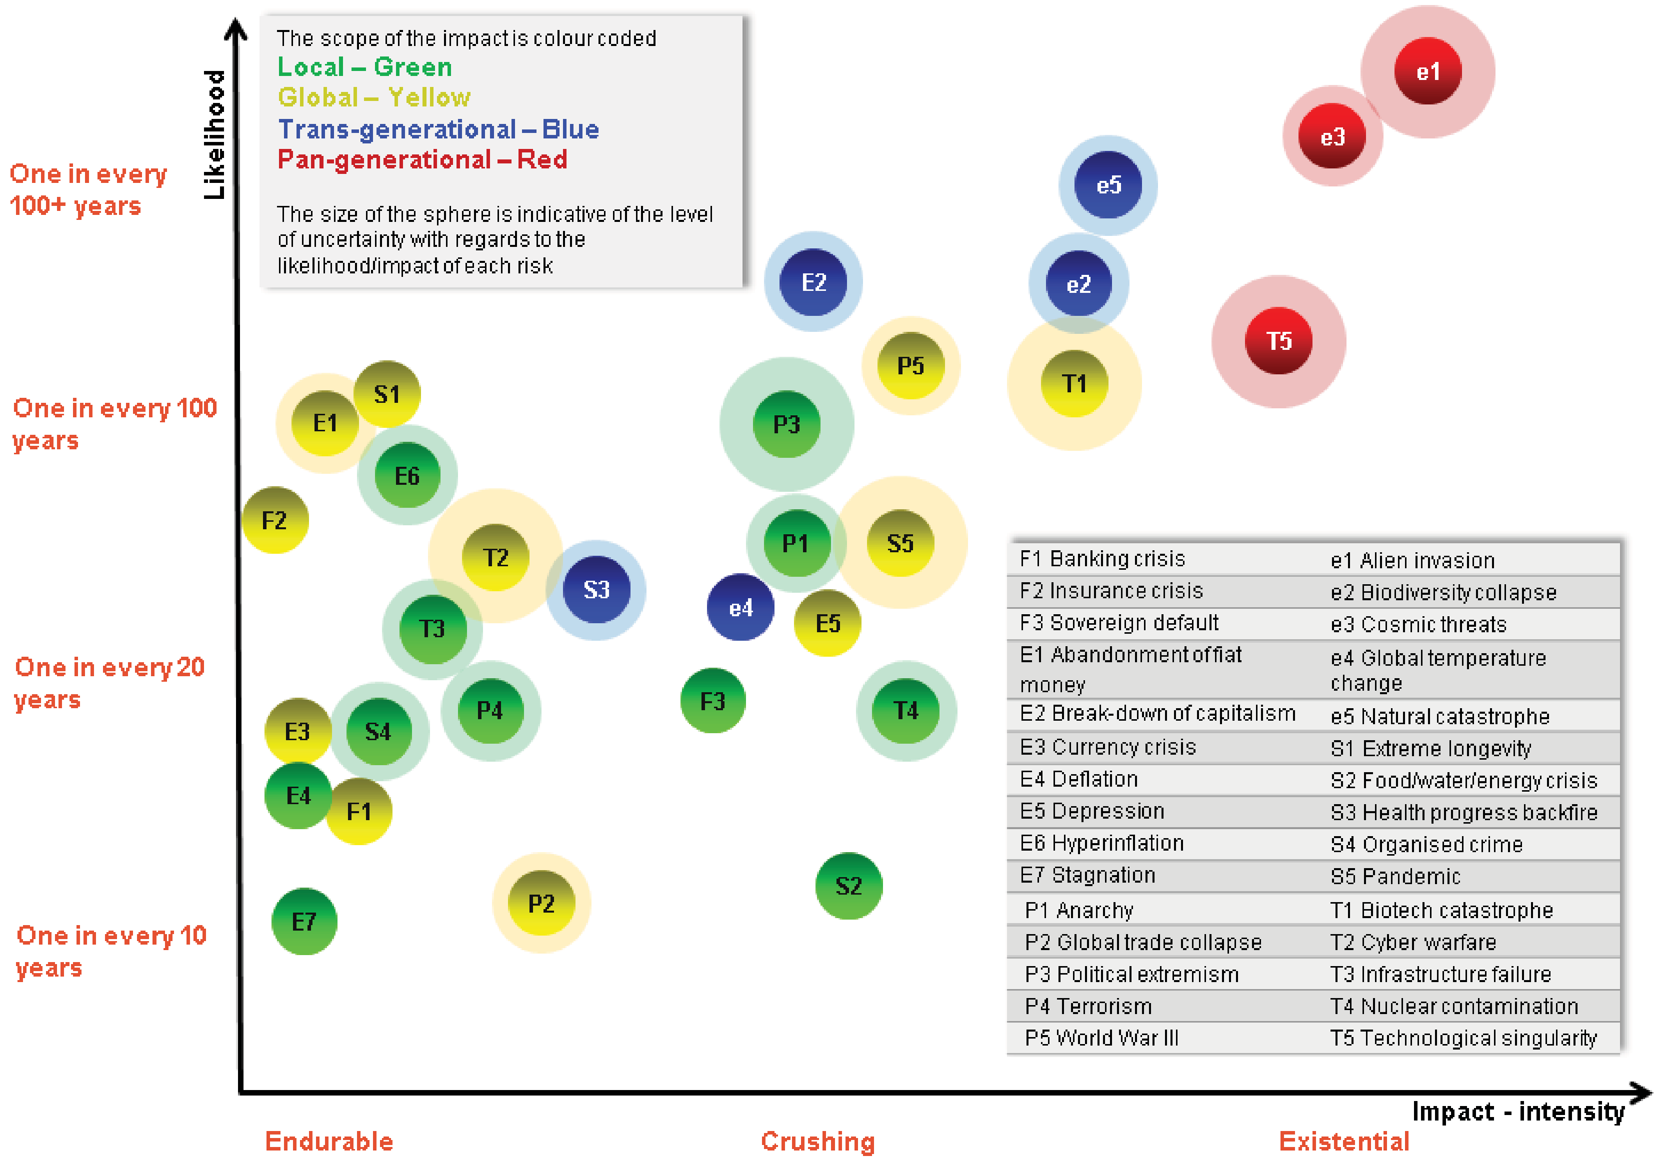

Figure 05 is perhaps more useful for assessing the results in aggregate. The good news is that extreme risks are much more likely to have temporal local or global effects (green and yellow) than lasting trans- or pan-generational effects (blue and red). While there is a general upwards slope towards top-right, implying that the worst risks were also the least likely, removing the trans- or pan-generational risks would make this relationship much less marked. The local (green) risks, whether endurable or crushing, as assessed as having likelihoods between 1-in-10 years and 1-in-100 years. In contrast the global risks (yellow) show more of an upward slope with crushing intensity only being associated with lower likelihood of occurrence.

Bringing this back to a retirement context we are interested in impacts on assets and liabilities. For ease of exposition, we can split **Figure 05** into three regions. The first comprises the six points in the top right corner, which are the existential and crushing-bordering-on-existential risks. Should these events happen we will not be particularly worried about the assets or the liabilities. The assets are likely to be worthless and the value of the liabilities may well have fallen to zero too. The second group comprises the 10 points to the left with crushing intensity of impact. For these events we expect the effect on assets to be global and materially negative. There would also be a reduction in the value of liabilities but as liabilities tend to be more local than assets the effect would be much more case specific. Somewhat crudely, if you are at war you lose lives, and your liabilities go down. If your neighbour is neutral, their liabilities don not change. On balance, therefore, we would expect asset losses to be larger than any reduction in liabilities and so funding levels would deteriorate.

The third group is the left-most fourteen points (endurable intensity). These are less homogenous but in general we would expect the impact on liabilities to be more muted. The two exceptions in this group are health progress backfire, which would reduce liabilities, and extreme longevity which would explicitly increase the liabilities. The effect on assets could span a massive range from a temporary drop in global asset values following an act of terrorism, to a permanent destruction of local asset values following a bout of hyperinflation.

Association. As already noted above, these risks are not entirely independent. We therefore also show an 'association' matrix in **Figure 06**. This is not a correlation matrix. Correlations require data to calculate and even then say nothing about causality. Instead we use the term association to communicate that this is a qualitative assessment of whether there is likely to be any causality between the events. If wanting to consider whether a particular event, X, might cause another, select X in the first column and read across the row. A blank cell means that, in our opinion, X does not cause Y to any material extent. An 'L' for low means that we believe X could cause Y, or is a contributory factor. An 'H' for high means we believe the causality is material, so X is likely to, or will, cause Y. For example, reading across the third row of entries shows that we believe a sovereign default (of a major country) (F3) *could* cause, or contribute to, a banking crisis (E3), the abandonment of fiat money (E1), a depression (E5), stagnation (E7), anarchy (P1), political extremism (P3), and World War III (P5). A banking crisis *is likely to* cause an insurance crisis (F2), and currency crisis (E3).

Figure 06 can also be read down the columns. In this case the column entries mean event Y could or is likely to 'be caused by' events corresponding to the cell entries. As an example, global temperature change (e4) is likely to be caused by biodiversity collapse (e2), cosmic threats (e3), and natural catastrophe (e5). Please note that this is an association matrix only between the risks we are considering, and therefore does not comment either way on whether there are other contributory factors outside our current consideration.

While many comments could be made about the associations, we make just a few. There is a significant clustering of associated risks within the financial, economic and political categories, which should not be surprising. Within these categories, an insurance crisis appears to be a relatively self-contained event in that it is assessed to be unlikely to trigger any of the other extreme risks considered here (the row F2 is empty). Similarly, terrorism (P4) is also relatively independent as both the row ('causes') and column ('caused by') have very few entries. It is worth noting that many of the risks in these categories are assessed as potentially causing both depression (E5) and stagnation (E7). These are both a negative shock to economic growth but are typically only distinguishable after the event. So while it would be possible for a depression to be followed by a decade or two of stagnation we would consider this a rare event and, rather, we would typically expect only one of these extreme risks to manifest.

Another high-density area of association is the environmental risks (rows e1 to e5) which appear powerfully causal so, contingent on one of them happening, we should expect other extreme risks to be triggered. Similarly, there is a small cluster at the intersection of politics and technology (rows P1 to P5 and columns T1 to T44), suggesting that a breakdown of political order is likely to, or could, trigger some technological risks. By contrast the social extreme risks (rows S1 to S5) have a very low causal power – they are unlikely to trigger further extreme events. However, the social risks could have multiple triggers as shown by the number of entries in columns S2 to S5. This is especially the case for food/water/energy crisis (S2) which is due to our assessment that the demand for water and food in particular could soon start to exceed current supply. With little current buffer it would not take much to tip the supply/demand balance the wrong way. Finally, we would note that there are some truly independent risks in our group. Columns e1, e3, S1, and T5 imply that an alien invasion, cosmic threats, extreme longevity and technological singularity are not triggered by any of the other extreme events considered here.

Figure 06. Extreme risk association matrix

	F1 Banking crisis	F2 Insurance crisis	F3 Sovereign default	E1 Abandonment of fiat money	E2 Break-down of capitalism	E3 Currency crisis	E4 Deflation	E5 Depression	E6 Hyperinflation	E7 Stagnation	P1 Anarchy	P2 Global trade collapse	P3 Political extremism	P4 Terrorism	P5 World War III	e1 Alien invasion	e2 Biodiversity collapse	e3 Cosmic threats	e4 Global temperature change	e5 Natural catastrophe	S1 Extreme longevity	S2 Food/water/energy crisis	S3 Health progress backfire	S4 Organised crime	S5 Pandemic	T1 Biotech catastrophe	T2 Cyber warfare	T3 Infrastructure failure	T4 Nuclear contamination	T5 Technological singularity	
F1 Banking crisis						L		L		L																					F1
F2 Insurance crisis																															F2
F3 Sovereign default	L	H		L		H		L		L	L		L		L																F3
E1 Abandonment of fiat money	L		L		L	L	L	L		L																					E1
E2 Break-down of capitalism	H	H	H			H	L	L	L	L	H	H	H									L		L				L			E2
E3 Currency crisis	L	L	L	L				L	L	L												L									E3
E4 Deflation			L					L		H																					E4
E5 Depression	H	L	H		L		H			L			L		L							L		L				L			E5
E6 Hyperinflation			L	H		H		L		L	L		H									L		L							E6
E7 Stagnation		L	L				H																								E7
P1 Anarchy			L			L		L		L		L	H		L							L		H		L		L	L		P1
P2 Global trade collapse								L	L	L	H		L		L							H		L			L	L			P2
P3 Political extremism			L	L	L	L		L	L	L		L		L	H							L					L		L		P3
P4 Terrorism											L		L													L		H	L		P4
P5 World War III	L	L	H	L	L	H						H	L				L					L			L	L	H	H	H		P5
e1 Alien invasion															L		L								L				L		e1
e2 Biodiversity collapse											H		H		L				H			H	H	L	L			L			e2
e3 Cosmic threats												H					H		H	L		H						H			e3
e4 Global temperature change										L	L		L				H					H	L		L			L			e4
e5 Natural catastrophe											H	H					H		H			H						H			e5
S1 Extreme longevity		H	L																			L									S1
S2 Food/water/energy crisis										L	H				L									L	L						S2
S3 Health progress backfire																									L						S3
S4 Organised crime										L	L																				S4
S5 Pandemic										L		L																			S5
T1 Biotech catastrophe											H	H										L	L		L						T1
T2 Cyber warfare	L														L													H	L		T2
T3 Infrastructure failure																						L									T3
T4 Nuclear contamination																						L						L			T3
T5 Technological singularity											L	L	L																		T5
	F1	F2	F3	E1	E2	E3	E4	E5	E6	E7	P1	P2	P3	P4	P5	e1	e2	e3	e4	e5	S1	S2	S3	S4	S5	T1	T2	T3	T4	T5	

Notes:
This table is a qualitative assessment of how likely the events in the left column can lead to events in the right columns (reading across the rows).
A blank cell means that X does not cause Y to any material extent
L (low) means that we believe X could cause Y, or is a contributory factor
H (high) means we believe the causality is material, so X is likely to, or will, cause Y
The table can also be read down the columns in a 'this risk is "caused by…"' sense. For example, E1 alien invasion is not caused by any of the other risks considered here
Source: Author.

Ranking. The final part of our assessment of these risks is to create a ranking of their importance. This introduces no new information but simply combines the four scores for each risk into a single number that can then be ranked. The intuition is straightforward. The more likely a risk, the higher up the ranking it should be. Likewise, the less uncertain a risk, the greater the intensity of impact and the larger the scope of the impact, the higher up the ranking a risk should be. The scoring system necessarily involves some element of 'rough justice' in that it will force a trade-off between, say, likelihood and intensity. How should a more-likely-but-endurable risk rank against a less-likely-but-crushing risk? Rather than worry too much about precise ranking, we think it is more useful to make an overall judgement as to whether the ranking is satisfactory. Are the 'right' sort of risks at the top, and are the bottom ones generally less important? Our ranking is shown in **Table 07**.

At the top of our ranking is food/water/energy crisis (S2). This is primarily driven by our assessment that this is one of the most likely risks and that there is *relatively* little uncertainty attached to either the likelihood or the consequences. The consequences themselves, locally crushing, are not particularly severe (again in relative terms) and so these are not a driver of the top ranking. The second ranked risk, stagnation (E7), differs only in respect of the intensity of impact which is assessed to be endurable rather crushing. In contrast the third ranked risk, global temperature change (e4), has much more severe consequences (trans-generationally crushing) but is assessed to have a lower likelihood of occurring.

It is interesting to note that the three extreme risks assessed as having the worst consequences (existential and pan-generational) all rank around or in the bottom third, with alien invasion (e1) ranked at 29. The poor consequences are 'diluted' by our assessment of low likelihood and relatively high uncertainty. Interestingly, the only one of these three risks we could do anything about, technological singularity (T5), ranks highest of the three at 19.

Right at the bottom is political extremism (P3). The locally-crushing consequences are not particularly severe, there is a high degree of uncertainty of how assets and liabilities would be affected, and the likelihood is relatively low all of which push it downwards in the ranking.

As alluded to above, different scoring systems are possible and they would produce somewhat different rankings. However the power of the ranking system is that it combines and trades-off the four risk scores in a consistent manner. Different weights could be applied, but the importance of a ranking system is to challenge preconceptions (and mitigate black swan biases). Whatever the weights, the ranking highlights the risks to prioritise when it comes to management actions. We will turn to this subject shortly, but before that we will consider the impact on retirement in more general terms.

Table 06 Extreme risk ranking

1.	S2 Food/water/energy crisis	16.	e2 Biodiversity collapse
2.	E7 Stagnation	17.	S5 Pandemic
3.	e4 Global temperature change	18.	P5 World War III
4.	E5 Depression	19.	T5 Technological singularity
5.	P2 Global trade collapse	20.	e3 Cosmic threats
6.	F1 Banking crisis	21.	T2 Cyber warfare
7.	F3 Sovereign default	22.	P1 Anarchy
8.	E3 Currency crisis	23.	E1 Abandonment of fiat money
9.	E4 Deflation	24.	S4 Organised crime
10.	S3 Health progress backfire	25.	E6 Hyperinflation
11.	T4 Nuclear contamination	26.	e5 Natural catastrophe
12.	S1 Extreme longevity	27.	E2 Break-down of capitalism
13.	F2 Insurance crisis	28.	T1 Biotech catastrophe
14	P4 Terrorism	29.	e1 Alien invasion
15.	T3 Infrastructure failure	30.	P3 Political extremism

Source: Author

Implications for retirement

What is retirement? In the introduction to this paper we drew a cartoon of human history suggesting that retirement used to involve raising grand children in return for food and board, and only recently has become self-sufficient leisure time. The point is that the definition of retirement is very flexible. It differs across history for an individual nation, it will differ between individuals of the same age in the same society, and it will differ between nations at any point in time. The key, we believe, is about expectations. How happy the generation that expected to perform unpaid childcare in retirement but got to live independently, albeit modestly. How distressed the generation that finds it cannot afford the annual Caribbean cruise it was expecting.

What are the explicit and implicit expectations for retirement currently? Restricting ourselves to the Anglo-Saxon world, we would characterise the expectations broadly as follows: to finish paid employment somewhere in one's 60's; to live self-sufficiently for another 20 or so years; to enjoy a standard of living equal to, or better than the previous generation; to benefit from political stability, clean air, running water, and health services; and to have access to some form of state insurance scheme to protect against poverty in old age. Clearly this list is incomplete, but we wanted to make the points that many expectations are implicit (such as political stability) and that adequate savings are necessary but not sufficient. For example, consider a shortage of running water. With large savings we can solve that problem by buying bottled water. But if bottled water is now very valuable, we become a target for crime, violent or otherwise. Retirement is as much about the continuation of the social order we expect as it is about having the money to buy 'stuff'.

Irrespective of any financial consequences that the extreme risks may have, they will pose a threat to the quality of retirement in terms of security and/or wellbeing – whether political, environmental, social or technological. We are not expecting prospective retirees to be overly concerned about an alien invasion (or anyone else for that matter), but the highest ranked risks would lower a retiree's quality of life relative to the previous generation, and likely in absolute terms too. A food/water/energy crisis, stagnation or depression would have a direct adverse effect. While global temperature change would have a direct impact it is likely to be more of a lottery of individual winners and losers (although negative in aggregate). The effects of a global trade collapse could be more indirect given transmission through various markets, but would nevertheless cause an adverse effect on purchasing power and/or the composition of the consumption basket.

Implications for retirement funds

While interesting in a general sense, we believe the consideration of extreme risks can be useful in helping to design more robust risk management processes. In our retirement context that means more robust investment portfolios, or more robust balance sheets. In essence the exercise of considering extreme risks is time spent on 'pre-mortems'. While a post-mortem seeks to establish the cause of death, pre mortems are about trying to determine in advance what could, colloquially, kill you. We believe that being adept at pre mortems means you are a better risk manager, and can react more flexibly in the event of an extreme event happening, particularly as the event is unlikely to evolve precisely along the lines predicted.

Consequently, the obvious application of extreme risk thinking is in stress-testing or scenario planning, but it is also constructive to consider whether the thinking can be incorporated within the process for managing a retirement fund's balance sheet.

One option would be to penalise the existing 'normal state' assumptions by slightly reducing expected returns, or pushing up volatilities, and/or correlations to reflect the impact of infrequent extreme events. A second option is dynamic switching of some sort. We either build two sets of assumptions ('normal' and 'extreme') or we design a second, extreme risk portfolio directly from first principles. Then 'all' that is left to do is successfully time the switch between the two, not forgetting the need to time the switch back so we can go on harvesting returns when the conditions are conducive.

The third option is a hedging overlay, and there are three broad alternatives available to us. The first is to hold cash; the second is to use derivatives; and the third is to hold a negatively-correlated asset. For a discussion of these alternatives please see Towers Watson 2011.

So how should retirement funds actually adapt in recognition of extreme risks? We would suggest a prioritisation exercise: first, worry about the events 'that can kill you', that is permanently impair the investor's mission. This should identify which extreme risks matter and which can be ignored. For the former, the right thing to do is to pay up for the insurance, given that the prioritisation exercise has shown the investor cannot afford to self-insure. Second, an investor should do the simple things. These would include ensuring the portfolio is as diversified across as many return drivers as possible; diversifying within asset classes; and creating a strategic allocation to cash to provide optionality. Finally, greater complexity can be added over time, assuming it passes a considered cost/benefit analysis. This is likely to involve adding long-dated derivative contracts in a contrarian manner, that is, when they are cheap rather than popular.

Extreme risks and the retirement anomaly

At the outset we posed the question of whether retirement in the modern sense is sustainable or a passing anomaly. Our review of extreme risks casts doubt on society's ability to defer sufficient current consumption to fund future consumption. A simple model of a retirement system where individuals work for 45 years, set aside 10% of earnings, and earn a real return of 3.5% each year, implies that they could live 21 years without working (Towers Watson 2012). But all the assumptions embedded in this model require the household to build up a pension or assets worth just over 10 times annual earnings, by the retirement date. If the population were distributed evenly by age, this would mean the economy would need to accumulate steady state pension wealth equal to 4.7 times total earnings, or around 235% of GDP.

Unfortunately, this target is far from met. The 13 largest pension markets in the world have pension assets amounting to less than 80% of GDP (Towers Watson 2013), implying too few resources available to support retirement. So we have either chronically under saved for retirement and need to do better, or other assets outside the pension system may be the answer. If the diagnosis is chronic under saving, it will be necessary to greatly boost saving without reducing the rate of return on investment. Unfortunately this could result in Keynes's paradox of thrift: even if it is rational for an individual to save more, if everyone saves more, the real rate of return is likely to be driven down. The same thought holds true for non-pension saving, and so it may not be possible for society to defer consumption in sufficient size to give people the retirement they currently expect.

It appears, then, that society will struggle to support a retired population in the style to which it aspires. The last half-century has witnessed a historical 'retirement sweet spot,' where rising expectations were affordable due to well-documented economic and social factors post-World War II, including favorable demographics and the run-up in debt levels. For the future, the outlook is less sanguine.

Conclusion

This paper has suggested that the world is a complex adaptive system. A consequence of this perspective is that the probabilities of extreme events are higher than anticipated by those accustomed to viewing the world linearly and predictably. We also noted that understanding that time is irreversible for decision makers implies that we must give greater weight to the consequences of outcomes and less weight to their likelihood. In turn, we must acknowledge that extreme risks can play havoc with the quality of life for retirees, and cause sustainability issues for retirement systems everywhere.

The range of potential consequences of these risks is very wide. Local endurable risks would be uncomfortable for retirement funds caught in the wrong locale, or with the wrong exposures, and would likely be enough to cause the weaker funds to become incapable of completing their mission. At the other end of the spectrum, existential and pan-generational risks represent a systemic and terminal outcome for retirement funds, and perhaps the human species. We also proposed a ranking system as a useful way to prioritize efforts to consider and manage potential risk exposures. Since political, environmental, social, and technological extreme risks are difficult to hedge, a relatively high cash weighting appears versatile and effective. At the edges some actions are possible, for example, the use of longevity swaps to hedge longevity risk looks very sensible for defined benefit funds (and would be useful for individuals with defined contribution funds if available in retail-friendly form). Beyond that a relatively high cash weighting appears the most versatile and effective hedge.

Given the hedging difficulties, the pragmatic solution is less about changing investment strategy and simply involves saving more and building a larger risk buffer. This course of action is currently being forced onto banks (higher capital buffers). While individually rational, this is likely to cause problems in aggregate. Consequently, the best outcome is likely to be a long, slow increase in risk buffers rather than a quick-fix solution.

There were issues that we did not have the space to discuss. Two are worth the briefest of mentions. Embedded within our likelihood scores are multiple stands of thought, but one of the most important is the absence of global governance, or 'G0[3]'. The political, social, technological, and two of the environmental risks are exacerbated by an absence of effective global governance. Conversely the development of such governance would, in our opinion, materially reduce the likelihood of these multiple risks. The second point is that we have identified several risks of 'technology running amok'. Clearly this is a risk, but does not have to happen. Technology has been, and will be, a substantial benefit to humankind. Nevertheless there remain important questions to be addressed concerning the rate of technological development relative to the level of our understanding, and our inability to predict the unintended consequences arising from complex adaptive systems.

In sum, extreme risks matter and they deserve more attention than given thus far. As a consequence, retirement for the masses is at serious risk, at least in terms of current expectations regarding length, quality of life, and degree of financial freedom. Alternatively, retirement as currently configured probably was never affordable, but this fact was obscured by demographic and debt trends over the past half century.

References

Adams, D. (1987). *Dirk Gently's Holistic Detective Agency,* UK, William Heinemann Ltd.

Adams, D. (1991). *The Long Dark Tea-Time of the Soul,* New York, Pocket Books.

Bostrom, N. (2013). 'Existential Risk Prevention as Global Priority,' *Global Policy*, 4(1):15-31.

Brock, H.W. (2003). 'Advent of a New "Investment Regime",' *SED Profile*, 68: I1-I25.

Caplan, B. (2006). 'The Totalitarian Threat,' in N. Bostrom and M.M. Cirkovic, *Global Catastrophic Risks*. Oxford, UK: Oxford University Press, pp. 504-519.

Gobel, C. and L.H.Wise (2012). 'Social Unrest in China,' *Europe China Research and Advice Network*.

Joy, B. (2000). 'Why the Future Doesn't Need Us,' Wired. April 2000: 8.04

Kilbourne, E.D. (2006). 'Plagues and pandemics: past, present, and future,' in N. Bostrom and M.M. Cirkovic, *Global Catastrophic Risks*. Oxford, UK: Oxford University Press: 287-307.

Kritzman, M. and Y. Li (2010). 'Skulls, Financial Turbulence, and Risk Management,' *Financial Analysts Journal*, 66(5): 30-41.

Larsen, P.T. (2007). 'Goldman Pays the Price of Being Big,' *Financial Times*. August 13.

Leitenberg, M. (2006). 'Deaths in Wars and Conflicts in the 20th Century,' *Cornell University Peace Studies Program Occasional Paper*, #29 3rd ed, ISSN 1075-4857

Lo, A. and M. Mueller (2010). 'WARNING!: Physics Envy May Be Hazardous To Your Wealth,' *Journal of Investment Management*, 8(2010): 13-63.

Lowrey, A. (2012). 'An Increase in Barriers to Trade is reported,' *The New York Times*. June 22.

Mandelbrot, M. and R.L. Hudson (2004). *The Misbehavior of Markets: A Fractal View of Financial Turbulence*. Basic Books.

Mitchell, M. (2009). *Complexity: a guided tour*. Oxford University Press.

Rees, M. (2003). *Our Final Century: The 50/50 Threat to Humanity's Survival: Will the Human Race Survive the Twenty-first Century?*, William Heinemann Ltd.

Soros, G. (2012). *Financial Turmoil in Europe and the United States*. PublicAffairs, pp. 44-45.

Sovacool, B.K. (2010). 'A Critical Evaluation of Nuclear Power and Renewable Electricity in Asia,' *Journal of Contemporary Asia*, 40(3): 369–400.

Towers Watson (2010). *Extreme risks*, TW-NA-2009-14398.

Towers Watson (2011). *Extreme risks – the 2011 update*, TW-2011-INV-00130.

Towers Watson (2012a). *The irreversibility of time*, TW-EU-2012-28847.

Towers Watson (2012b). *The impossibility of pensions*, TW-EU-2012-25855.

Towers Watson (2013). *Global Pension Asset Study 2013*.

United States National Intelligence Council (NIC) and European Union Institution for Security Studies (EUISS) (2010). *Global Governance 2025: at a Critical Juncture*. NIC 2010-08.

West, G. (2011). ‘Emergence of “Universal” Time in Living Systems from Cells to Cities,’ *Santa Fe Institute.*

World Economic Forum (WEF) (2013). *Global Risks 2013 – An initiative of the Risk Response Network.*

Zielinski, G. A.; P.A. Mayewski, L.D. Meeker, S. Whitlow, M. Twickler and K. Taylor (1996). ‘Potential Atmospheric impact of the Toba mega-eruption ~71,000 years ago’, *Geophysical Research Letters* (United States: American Geophysical Union) 23 (8): 837–840.

Extreme risk appendix

The extreme risks developed in the text are described in greater detail here.

Financial risk.

F1 Banking crisis. In the next crisis, central banks might be unwilling or unable to supply liquidity, and governments may be unwilling or unable to supply capital, causing banking and real economic activity to stop. Developments that could threaten future bank solvency could include (1) a continued drop in real estate prices, (2) increased corporate defaults and (3) poor economic conditions in general. What normally follows a banking crisis includes seizure within the financial markets triggering a flight to safety, or a swift transmission to the real economy threatening trade and businesses. Economic growth halts as the availability of credit evaporates with the possibility of triggering a depression. Historically banking crises have seen the nationalisation of one or more banks, so a future issue would be the state of government finances at that point in time, and whether this action was possible without triggering one of the other (economic) risks (for example, sovereign default) identified in this paper.

F2 Insurance crisis. This extreme risk is the insolvency of the insurance industry, that is, the assets are (become) smaller than the liabilities. This could occur through inadequate capitalisation of risky asset positions (the value of risky assets fall more than anticipated and do not recover) or through 'basis' risk (the safe assets held against the liabilities underperform, for example through sovereign default). We would expect a similar unfolding of events to that of the banking crisis, namely 'strong' insurers would be encouraged to take over failing insurers. Failing that, governments would provide some form of backstop (if able to). We assume that any such crisis would affect life insurance rather than general insurance (which may only be adversely affected by sovereign default given the nature of the assets held). The main consequence would therefore be lower receipts by the private sector, and hence lower consumption.

F3 Sovereign default. This extreme risk refers to the non-payment by a major sovereign borrower, causing market panic and adversely disrupting the global economy. Sovereign default ('restructuring' is the more polite euphemism) has also been surprisingly common, and not just among emerging countries. It is easily possible to document over 140 sovereign defaults over the centuries (less strict definitions would yield many more) – and that number includes the likes of Germany, France, Italy and Japan in the 20th century. There is clearly a benefit to defaulting – not paying back what you owe allows you to consume more immediately (and possibly into the future). The costs associated with sovereign default fall into two broad types – penalty costs and output costs. Penalties are applied by external creditors on the cost or ability of defaulters to access future finance. Output costs refer to the drop in production and therefore consumption of the defaulting nation. As domestic banks tend to be large holders of sovereign debt, particularly just before a default, the act of default can trigger a banking crisis and therefore an economic crisis. There has been much alarmist talk about the size of budget deficits and hence the required size of sovereign debt issuance. We believe sovereign default is more likely if economic growth remains stuck at a low level; if taxes are not or cannot be increased; and if governments do not (or cannot) reduce spending. The likelihood also increases when the tax burden falls on workers who do not believe it is their issue (the next generation). Sovereign default is usually preceded by several actions: regulatory changes (wide range of possibilities from changing state pension benefits, tax rates, permissible investments, and so on); covert sequestration (for example tax breaks for holding government bonds) to compulsory sequestration (forced conversion of cash to government bonds, making the holding of gold illegal, and so on).

Economic risk.

E1 Abandonment of fiat money. The fiat money system, exclusively used in every major economy nowadays, is the system where a government agrees to accept non-convertible paper money in payment of taxes and debts. Fiat money is declared by a government to be legal tender (the term derives from Latin and means 'let it be done'), and therefore does not have any intrinsic value. Its value is dependent on the relative scarcity and the degree to which people trust it. A gold standard (or gold-exchange standard) is associated with a government guarantee that paper notes are freely convertible into pre-set, fixed quantities of gold. Since the collapse of the Bretton Woods system in 1971, we have now been living in a fiat-money world for over 40 years. From a historical point of view, however, a repeated shifting between fiat and gold standards is normal. The United States has so far avoided hyperinflation by shifting between a fiat and gold standard over the past 200 years.[4] In a fiat money system, central banks have no upper limit to an expansion of the money supply while in a gold standard monetary system the supply of money is limited by the relatively stable (in the long term) pace of gold mining, thus limiting central banks' efforts to significantly stimulate the economies. The choice is then between a stable economy but unstable monetary policy (fiat money) and a stable monetary policy but unstable economy (gold standard). The escalating money creation around the globe has prevented another depression-like scenario since the global financial crisis. All of these efforts would become impossible under a gold standard structure. Nonetheless, should these efforts result in rising inflation or even hyperinflation the return to a gold standard would have a higher probability in future. A deepening distrust of fiat money among investors would have long-term investment implications, for example an intensified search for alternative assets as a store of value.

Before gold ever came back as a 'standard' it would be likely to benefit from any increasing distrust of central banks and the paper money they issue.

E2 Break-down of capitalism. The headline here is that capitalism is not in crisis but 'Capitalism IS the Crisis'.[5] It is the ultimate economic extreme risk where distrust in the private capital/property system causes a collapse in economic activity and asset prices. Capitalism's basic premise is that the pursuit of self-interest and the right to own private property are morally defensible and legally legitimate. In a pure capitalist economy, the market drives the allocation of resources and any economic decisions. In contrast, socialism advocates public ownership, in which governments determine the means of production and the allocation of wealth. Arguably, however, a fully market-driven economy (that is, the capitalism in the textbook) has never existed, neither has a completely centralised economy. In our view, the most likely scenario is moving along a spectrum where at one end the market is king (minimum regulation) towards the other end where we could see more onerous regulations and government intervention and control of the economy. The extreme risk, however, is the demise of the capitalist system and the end of the market as the primary means of resource allocation. In this scenario, the economy is likely to run a higher risk of failure. Low productivity will result in sluggish economic growth in the long run. The link between productivity and reward is lost. In addition, centralised power increases the problem of corruption (chance and/or scale) which inflicts substantial economic costs. As governments take on allocating resources, the (private) investment activities will collapse or even be terminated. The economy is likely to be associated with extreme uncertainty and a large amount of wealth destruction during the transition period. On contemplating the break-down of capitalism, investors should probably worry more about the return *of* their investments instead of the return on their investments.

E3 Currency crisis. 'Currency crisis' is an alternative term for 'balance of payment crisis' and is therefore, technically, the breaking of a fixed exchange rate. In a looser sense, it can also mean an expectation of a significant self-fulfilling devaluation of a major currency. Ideally economic management is used to maintain balance. This could be through policies to make domestic business more efficient (raising exports) or changing interest rates to attract/deter capital inflows. Therefore the movement in the exchange rate can be thought of as a safety valve that had to blow because other (painful) economic adjustments were not made – for example, raising interest rates and/or taxes. For a fixed exchange rate, this will occur when the central bank runs out of reserves and can no longer defend the exchange rate. A currency collapse severely reduces a country's purchasing power and hence wealth. To the extent that domestic borrowing has occurred in foreign currencies, the cost of servicing that debt will rise dramatically, and hence immediately increase the risk of default. The direct impact on asset returns is through the currency – domestic investment in domestic assets will be unaffected; domestic investment in foreign assets will benefit substantially if unhedged; investment by foreigners in domestic assets, if unhedged, will suffer very poor returns. The indirect effects are more complicated as the crisis will only have occurred because of some underlying economic imbalance.

E4 Deflation. This refers to goods and services prices falling for a long period (the inflation rate falls below 0%) in a major economy. Historically not all episodes of deflation correspond with periods of poor economic growth, particularly when deflation was caused by technological progress that created significant economic growth. The extreme risk occurs when deflation aggravates recessions and leads to a deflationary spiral. A deflationary spiral is a situation where decreases in price lead to lower production, which in turn leads to lower wages and demand, which leads to further decreases in price – a vicious circle, where a problem exacerbates its own cause. Deflation increases the real value of debt, causing a transfer of wealth from borrowers to the benefit of savers. Confused pricing signals cause under-consumption and under-investment at the cost of jobs and future economic growth. This is often noted as a liquidity trap by Keynesian economics in which people hoard cash because they expect deflation, and central banks are incapable of stimulating the economy by lowering interest rates (in a liquidity trap short-term interest rates are typically near zero). Long periods of persistent deflation can be severely detrimental to economic growth: rising real interest rates, drastic declines in output and persistently high unemployment. The Great Depression was regarded by some as a deflationary spiral. Deflation should result in slower growth and rising real debt service costs, which will depress corporate earnings growth and equity returns. The increase in real discount rates will have a further negative impact on equity prices.

E5 Depression. The extreme risk is a deep trough in economic output, massive increase in unemployment, restriction of credit, and shrinking investment in a major economy. The current risk of depression in the West appears to have been reduced through policy action, but remains an extreme risk – in that it may not be possible for governments to counteract any future drop in demand, should that occur. There has been an extended period of over-consumption (by Western consumers) meaning that businesses have built productive capacity to satisfy a level of demand that is unlikely to be reached for a number of years, as Western households increase their savings rate. The primary consequence of a depression is typically a sharp and prolonged increase in unemployment. The depth of the trough means that a long period of recovery is required before there is pressure to hire new workers. The subsidiary effects are therefore a drop in consumption, restriction of credit, shrinking output and investment, and numerous bankruptcies. Depressions can trigger deflation or hyperinflation, adding further complications. Excessive leverage in the system can interact with depression – a self-reinforcing fall in asset values can cause further defaults, bankruptcies, falling incomes and rising unemployment, causing or prolonging economic depression.

E6 Hyperinflation. The extreme risk is inflation being very high or 'out of control' in a major economy, a condition in which prices increase rapidly as money loses its value. Definitions used by the media vary from a cumulative inflation rate over three years approaching 100% to 'inflation exceeding 50% a month'. As a rule of thumb, hyperinflation is often reported for short intervals, often per month. The traditional quantity-theory-of-money view tends to attribute hyperinflation to unchecked budget deficits, leading to a rise in the supply of money and consequently higher prices, while James Montier of GMO[6] provides an alternative view that money supply is endogenous and that budget deficits are often caused by hyperinflation rather than being the source. He notes that hyperinflation shares several common characteristics such as large supply shocks, big debts denominated in a foreign currency and distributive conflict/transmission mechanism. Regardless of the debate about the root causes of hyperinflation, it is usually accompanied by a widespread unwillingness to hold the money for more than the time needed to trade it for something tangible to avoid further loss. Hyperinflation wipes out the purchasing power of savings, provokes extreme consumption and hoarding of real assets, causes the monetary base to flee the country, and investment ceases. Historically, there have been numerous episodes of hyperinflation in various countries, followed by a return to 'hard money' (some form of non-devaluing medium of exchange). Hyperinflation is often associated with wars (or their aftermath), economic depressions and political or social upheavals. The general population prefers to keep its wealth in non-monetary assets or in a relatively stable foreign currency. Amounts of local currency held are immediately invested to maintain purchasing power. Prices may be quoted in a foreign currency.

E7 Stagnation. Compared to depression which we define as having a very painful contraction phase but then followed by faster recovery, stagnation is a prolonged period of little or no economic growth (in a major economy), usually accompanied by high unemployment. Economic stagnation can be caused by excess leverage build-up in the system and the subsequent prolonged period of deleveraging. It could also be the result of a catastrophe or demographic changes. One of the latest and widely researched economic stagnations is the so called 'The Lost Decade(s)' of Japan since the early 1990s when its massive scale speculative asset price bubble collapsed in a catastrophic manner, aggravated by a declining and aging population. Japan's economy has stagnated for more than two decades and its real GDP growth rate from 1990 has been less than 1% per year, noticeably less than the trend growth rates achieved over the decades before the stagnation. During the same period, Japan has also experienced a secular decline in the employment-to-population ratio and a secular increase in the unemployment rate. Persistent deflation has become a norm. Real earnings declined steadily which limited the growth in real private consumption. Economic stagnation is normally associated with low real interest rates, under which conditions nominal bonds tend to outperform while risky assets such as equities suffer.

Political Risk.

P1 Anarchy. The Arab Spring that started in December 2010 has removed existing rulers in Tunisia, Egypt, Libya and Yemen to date. It was prompted by dissatisfaction with the rule of governments and likely by wide gaps in income levels. Welfare cuts and unemployment during the global financial crisis fuelled protests and anxiety across Europe. According to

Europe China Research and Advice Network, social unrest in China has been increasing at an alarming rate – 8,700 'mass incidents' were recorded in 1993 and by 2005 it had grown tenfold to 87,000 and estimates for 2010 range between 180,000 and 230,000, highlighting an increasing threat to the stability of world's second largest economy (Gobel and Ong 2012). In a world of growing income inequality and hyper-connected communication, the risk is an extreme form of social disorder in a major state, resulting in a loss of power by government and causing its economy to collapse.

P2 Global trade collapse. Protectionism is the policy of restricting trade with the aim of 'protecting' businesses and workers in the domestic economy from the full force of external competition. There have been a number of studies that suggest an increase in barriers to trade since the global financial crisis (Lowrey 2012). The concern is that short-term political expediency can override long-term economic logic with the extreme risk being a populist backlash against cross-border mobility of labour, goods, and capital, causing global trade and investment to collapse. The consequence will include more uncertainty in financial markets, greater fragmentation of capital markets and eventually a reversal in globalization.

P3 Political extremism. During the twentieth century, many nations suffered under extraordinarily brutal governments, which intended to hold total authority over the society and seek to control all aspects of public and private life (totalitarianism). The Soviet Union and Nazi Germany are the two most-studied totalitarian regimes. The risk of political extremism is defined by the rise to power in a major economy of an oppressive government (including but not limited to totalitarianism). Political extremism typically causes a large number of civilian deaths (by modern calculations, the Soviets killed approximately 20 million civilians, the Nazis 25 million) and could become a major threat to global peace (Nazi Germany directly caused World War II). Bryan Caplan from George Mason University speculates that the chance of a world totalitarian government emerging during the next 1,000 years and lasting for 1,000 years or more is 5% (Caplan 2006).

P4 Terrorism. 9/11 caused almost 3,000 deaths and the Dow Jones Industrial Average index to fall by more than 14% within the first week of market reopening. New York City's GDP was estimated to have lost US$27.3 billion for 2001-2012.[7] Its impact extended beyond geopolitics into society and culture in general. The extreme risk here is a major ideologically-driven terrorist attack of a similar, or larger, scale to 9/11, targeted at a region of global economic and/or political importance and inflicting large-scale human and financial damage.

P5 World War III. This extreme risk is a military war involving many of the world's most powerful and populous countries causing multiple-millions of deaths. One consequence of war is the destruction of capital – both physical and human capital. War tends to kill those in prime ages (predominantly males), which leaves a reduced younger workforce base and in turn reduces economic output and consumption. World War II caused deaths of between 65 and 75 million, and the total number of deaths in wars and conflicts for the entire 20th century was between 136.5 and 148.5 million (Leitenberg 2006). The availability of weapons of mass destruction means the next world war could destroy an order of magnitude more capital than the previous ones. As Albert Einstein put it 'I know not with what weapons World War III will be fought, but World War IV will be fought with sticks and stones'. The invention of nuclear and biological weapons raises the possibility that the future war could put much of the human race at risk.

Environmental risk

E1 Alien invasion. An alien invasion is a very common theme in science fiction stories and films despite the fact that evidence of extra-terrestrial life has never been documented. NASA's Kepler mission to identify earth-size planets around stars was launched in March 2009 and has already discovered thousands of candidates (2,740 planet candidates and 105 confirmed planets as of 11 Feb 2013[8]), including one the size of earth. The range of outcomes of an alien life contact can be vast and entirely unpredictable but if the contact is indeed hostile it is more likely that human race is unable to defend itself due to the potentially overwhelming technological gap. The extreme risk is therefore an invasion of non-peace-seeking extra-terrestrials that look to either remove the planet's resources or enslave or exterminate human life.

E2 Biodiversity collapse. It is estimated that less than 1% of the species that have existed on earth are extant and there have been five known mass extinctions since life began on earth that led to large and sudden drops in biodiversity.[9] Human activity has accelerated the species loss and these losses could reach a point beyond which it becomes irreversible. It is believed earth is not far away from its sixth mass extinction. Although about 80% of humans' food supply comes from just 20 types of plant, humans use at least 40,000 species. Earth's surviving biodiversity provides resources for increasing the range of food and other products suitable for human use, although the present extinction rate shrinks that potential. The subsequent destruction of the world's ecosystem can cause the loss to humans of ecosystem services: provision (food and clean water), regulation (climate and disease), support (nutrient cycles and crop pollination) and culture (spiritual and recreational benefits).

E3 Cosmic threats. There are risks arising beyond earth, such as a major meteorite impact, being pulled out of orbit (or the solar system) by a passing asteroid, or a giant solar flare (the effects of which would be compounded if during a reversal of the earth's magnetic field). The impact of these events could range from severely inconvenient to existential. A big enough solar eruption could trigger a magnetic storm and damage electricity distribution lines or disable critical communication and navigation systems, while a 10 kilometer wide meteorite (like the one that hit earth around 65 million years ago causing, as widely believed, the extinction of dinosaurs) could release 100 million megatons equivalent of energy. It is estimated that such a meteorite could trigger magnitude 10 earthquakes and a 300-metre high tsunami spreading to all of the earth's coastal regions, costing millions if not billions of human lives. Noxious gases and dust would then accumulate in the atmosphere cutting out sunlight and potentially terminating all lives that survived the direct impact – a mass extinction event.

E4 Global temperature change. There is little doubt in science that rising greenhouse gas emissions produced by human activities are leading to rising global temperature. Natural feedbacks (for example, the ice-albedo feedback means that melting ice reveals darker land and water surfaces below, which absorb more solar heat, causing more melting and warming) in the system have the potential of amplifying global warming. It is expected to be followed by serious consequences including extreme weather being more frequent, and rising sea levels (of several meters) making much of the current coastal communities uninhabitable. The extreme risk is earth's atmosphere passes a point of no return and tips into a less-habitable state. On the other hand, while gaining less support in the science community, earth's surface and atmosphere could experience excessive cold slipping into an ice age. This could be caused by a drop in the sun's emission of energy (for a temporary but prolonged period), or by another extreme event such as a meteorite strike or supervolcano. In either situation, habitable areas will be significantly reduced, causing large scale migration and reducing the quality of life for most of humankind.

E5 Natural catastrophe. These are the disasters resulting from natural processes of the earth including earthquakes, tsunamis, hurricanes, flooding (including atmospheric river storms) and volcanic eruptions. The extreme risk would either be a confluence of connected extreme natural catastrophes (for example, a magnitude 10 earthquake, causing a giant tsunami and triggering volcanic eruptions) or the eruption of a supervolcano. The latter would cause global effects on climate from the ash fallout and aerosol clouds ('volcanic winter'), agriculture (collapse as a result of the loss of one or more growing seasons), health (famine and spread of infectious disease), and transportation (air travel halted for years). It is believed that a supervolcanic event at Lake Toba around 71,000 years ago led directly to a cooling event that lasted over 1,000 years (Zielinski et al 1996).

Social risk

S1 Extreme longevity. A major breakthrough in medical or human genome science – it is hoped that the cure for common condition such as heart disease, cancer and stroke may be in the offing – could result in an unanticipated, significant increase in life expectancy for many, or the majority, of humans. A direct impact of people living longer on a retirement plan is increased liabilities. In addition, even though life expectancy has increased steadily in recent history these gains do not necessarily lead to better health in later life. The risk therefore also includes an emergence of a society of a growing number of the elderly who suffer chronic but nonfatal diseases – people live longer but their 'productive' years stay more or less the same. The economy will be struggling to support the health care of a mass of the elderly who are in need of long-term health care.

S2 Food/water/energy crisis. It was estimated in 2010 that 600 million people in 21 countries were facing either cropland or fresh water scarcity, and that number is projected increase to 1.4 billion people in 36 countries by 2025. Over one billion people live in areas where human use of available water supplies exceeds sustainable limits and by 2025 this figure is projected to rise to 1.8 billion, with up to two-thirds of the world's population living in water stressed conditions (NIC and EUISS 2010). On the energy side, the supply of fossil fuels has a known limited time span while no viable alternatives are currently available with comparable energy returns on energy invested (EROEI). There is a risk that the necessary technological breakthrough will not arrive in time to prevent a global economic collapse due to an energy crisis. Consequently, given the current fine balance between supply and demand

and the projections of demand growing faster than supply for food, water, and energy, we see this as a particular area of vulnerability. The extreme risk refers to the occurrence of a major shortfall in the supply of, or access to, food/water/energy for a large proportion of the world's population, causing severe societal issues such as widespread death and damage to the quality of life for many survivors.

S3 Health progress backfire. Modern medicine has been consistently meeting existing and new diseases with new treatments, giving rise to improved human health. There is no guarantee that the rate of medical advancement can always outpace the rate of pathogen evolution and a catastrophic event could emerge should biological mutation eventually outpace human innovation. This could result from the unintended consequences of current healthcare practices such as antibiotic resistance. The World Economic Forum (WEF 2013) warns that we are decades behind in comparison with the historical rate at which we have discovered and developed new antibiotics and none of the drugs currently in the development pipeline would be effective against certain killer bacteria. Social trends such as widespread mental health problems and obesity are additive to the problem. Stephen Petranek, then editor-in-chief of Discover magazine, points out in a TED talk that despite improved physical health, the human race is mentally falling apart – one in five people in the West is believed to be clinically depressed.[10] The extreme risk from a societal point of view is a massive increase in morbidity for a large proportion of the population. Not only does this directly reduce quality of life, but would also reduce economic output. From a retirement viewpoint the extreme risk is that the increase in morbidity is not accompanied by a reduction in longevity. In other words, economic output falls and liabilities increase.

S4 Organized crime. Organized crime is a common reality for most if not all countries. The UK Home Office suggests that organized crime costs the UK between £20 and £40 billion each year[11] and its impact is felt by the state, businesses, communities, families and individuals. The extreme risk is a significant increase in the scale of illegal operation in a major economy to the extent that legitimate economic activity becomes non-viable. Extreme forms of organized crime could bring severe disruptions to normal activities in affected areas, typically associated with high homicide rates, wide use of illegal drugs and the collapse of legal business activity potentially followed by social unrest.

S5 Pandemic. Recent pandemics (for example, SARS, avian flu and swine flu), despite being successfully contained (for now?), demonstrate how easily deadly viruses can mutate and history is full of significant pandemics with an extremely high number of causalities. For example, it is believed that Plague of Justinian in AD 541-542 killed 50% of the world's population; the Black Death in the 13th century caused the death of one-third of the population of Europe, and 'Spanish flu' during 1918-19 killed 20-50 million people (Kilbourne 2006). [We need to distinguish between those pandemics occurring before the advent of modern medicine and those after. For example, the Black Death is believed to have been a bacterial infection which would, today, be treated with antibiotics. However, please note the threat of antibiotic resistance referred to in S3.] Pandemics can be attended by high morbidity within a very short period of time (for example, influenza), increasing the difficulty for effective vaccines to develop in time. Modern travelling patterns make it almost impossible to contain a contagious disease within a specific region. While we have relatively more knowledge about human disease pandemics than other events, the extreme risk here is a pandemic of a new highly infectious and fatal disease that spreads through human, animal or plant populations worldwide.

Technological risk

T1 Biotech catastrophe. DNA sequencing and synthesizing machines are available to anyone with enough money to afford a used car. Nasty nucleotide sequences such as the Ebola virus and the 1918 influenza virus are accessible online and genetic engineering of viruses is much less complex and far less expensive than sequencing human DNA. This makes it a lot easier to apply this technology to destructive uses than constructive ones. Adding to the problem is the fact that the biotech industry is highly unregulated. Regulating and controlling current and new developments would require strong global governance which the world currently lacks. British cosmologist and astronomer Martin Rees speculates that by the year 2020, an instance of bio-error or bio-terror will have killed a million people (Rees 2003), which is the extreme risk considered here.

T2 Cyber warfare. This refers to politically-sponsored computer hacking to conduct sabotage and espionage on a national or global-power scale. It is reported that a series of cyber-attacks on businesses and institutions in the United States have prompted fears of a looming 'cyber cold war' and outgoing United States Defense Secretary Leon Panetta predicted a cyber-version of Pearl Harbor might soon take the United States by surprise. Cyber war could cause severe damage to physical infrastructure – bridges, tunnels, air traffic control, and energy pipelines. Social security, financial and medical systems connected to the internet could all become the target of cyber-attacks. A cyber-attack on the defense system might precede a military attack in all future wars.

T3 Infrastructure failure. This extreme risk refers to an interruption (prolonged but not permanent) of a major infrastructure network due to either human activity (for example, cyber-attack), natural disasters (for example, earthquake or flooding), or even cosmic threats (for example, giant solar flare). An extended shut-down of a critical network or electricity grid would bring increasing disruption to economies within the geographical area of impact. People's basic needs would be threatened in such circumstances raising the possibility of social unrest and law-breaking behaviors for survival.

T4 Nuclear contamination. The risk is a major nuclear accident or attack that leads to lethal effects on individuals and large radioactivity release to the environment. It is reported that worldwide there have been 99 accidents at nuclear power plants from 1952 to 2009 (defined as incidents that either resulted in the loss of human life or more than US$50,000 of property damage), totaling US$20.5 billion in property damages (Sovacool 2010). One of the worst nuclear contamination events to date is the Chernobyl disaster which occurred in 1986 in Ukraine, killing 30 people directly, causing thousands of indirect deaths due to radiation-induced cancer, as well as damaging approximately US$7 billion of property.[12] Nuclear contamination can also be a direct consequence of a nuclear terrorist attack or a full-blown nuclear war among states.

T5 Technological singularity. This refers to an extreme risk resulting from technological advancement proceeding beyond the point of human understanding. It is possible that the creation of a computer more powerful than the human brain, which can then design and build an even more advanced machine, would create an environment where human survival is at risk. Bill Joy, then Chief Scientist at Sun Microsystems, has argued that '21st century technologies – robotics, genetic engineering, and nanotech – are threatening to make humans an endangered species' because 'they share a dangerous amplifying factor: they can self-replicate' (Joy 2000). Another possibility is a so-called nanotechnology 'grey goo' scenario, in which nano-robots self-replicate in an uncontrolled manner and eventually consume everything on the earth.[13] The University of Cambridge has recently established a research Centre named 'The Centre for the Study of Existential Risk', devoted to studying possible catastrophic threats posed by present or future technology.

---

[1] See the Santa Fe Institute's website http://www.santafe.edu/about/history/

[2] With extreme risks all probabilities are qualitative, however derived. By definition we are dealing with very small sample sizes of previous occurrences so even using extreme value theory to fit a tail distribution requires the application of significant qualitative judgement

[3] A term and concept coined by Ian Bremmer of Eurasia Group

[4] http://www.kwaves.com/fiat.htm

[5] If you are interested in this subject, there is a long documentary that explains the capitalist crisis (running time 100 minutes): http://topdocumentaryfilms.com/capitalism-is-the-crisis/

[6] "Hyperinflations, Hysteria, and False Memories", James Montier, GMO white paper, February 2013

[7] Sourced from Wikipedia at http://en.wikipedia.org/wiki/September_11_attacks

[8] http://kepler.nasa.gov/

[9] All data in this section from http://en.wikipedia.org/wiki/Biodiversity

[10] The presentation can be watched at http://www.ted.com/talks/stephen_petranek_counts_down_to_armageddon.html

[11] UK home office data at http://www.homeoffice.gov.uk/crime/organised-crime-strategy/

[12] Sourced from Wikipedia at http://en.wikipedia.org/wiki/Nuclear_and_radiation_accidents

[13] Sourced from Wikipedia at http://en.wikipedia.org/wiki/Grey_goo

Thinking Ahead

This publication is written by members of our Thinking Ahead Group (TAG) who are part of Investment at Towers Watson. Their role is to identify and develop new investment thinking and opportunities not naturally covered under mainstream research. They seek to encourage new ways of seeing the investment environment in ways that add value to our clients. The contents of individual articles are therefore more likely to be the opinions of the respective author(s) rather than necessarily representing the formal view of the firm. No action should be taken on the basis of any article without seeking specific advice. If you would like to discuss any of the areas covered in more detail, please get in touch with the consultant who normally advises you at Towers Watson, or:

Tim Hodgson
+44 1737 284822
tim.hodgson@towerswatson.com

Liang Yin
+44 1737 284762
liang.yin@towerswatson.com

About Towers Watson

Towers Watson is a leading global professional services company that helps organisations improve performance through effective people, risk and financial management. With 14,000 associates around the world, we offer solutions in the areas of benefits, talent management, rewards, and risk and capital management.

Towers Watson
21 Tothill Street
Westminster
London
SW1H 9LL

towerswatson.com

TOWERS WATSON