

Extreme Risks and the Retirement Anomaly

Tim Hodgson

October 2013

PRC WP2013-28

Pension Research Council Working Paper

Pension Research Council

The Wharton School, University of Pennsylvania

3620 Locust Walk, 3000 SH-DH

Philadelphia, PA 19104-6302

Tel: 215.898.7620 Fax: 215.573.3418

Email: prc@wharton.upenn.edu

<http://www.pensionresearchcouncil.org>

All opinions, errors, findings, interpretations, and conclusions of this paper represent the views of the authors and not those of the Wharton School or the Pension Research Council. © 2013 Pension Research Council of the Wharton School of the University of Pennsylvania. All rights reserved.

Extreme Risks and the Retirement Anomaly

Abstract

Extreme risks are potential events that are very unlikely to occur but which would have a significant impact should they happen. The global financial crisis and its aftermath have demonstrated that risk management cannot afford to stop at the 95th percentile (VaR95) and a holistic risk management framework should include very unlikely, but potentially high impact, events. In essence we are calling for an ‘overweighting’, whether of attention or resources, to be applied to the very low-probability, but bad, outcomes. Why? Because we only live once. Much of finance and economics assume we have infinite lives all running in parallel. This means that the impact of extreme risks can be masked in expected return calculations by frequent positive outcomes. For an individual planning for retirement, a one-in-a-million chance matters because they live in a single universe and face problems in series, not parallel. If that individual gets the one-in-a-million path it could well be terminal – and how they would have fared on the other 999,999 paths is of no useful interest.

This paper considers a number of extreme risks. The focus is directed more towards non-economic risks on the grounds that these have received less attention than the more obvious financial and economic risks. However there is one economic risk that we do not think is getting enough attention, namely whether it is possible for society to provide the retirements to which its citizens aspire. The current consensus, among practitioners if not the general population, is that people must save more for their retirement unless they are willing to work until they are very old. While a sensible course of action for an individual, can this be done in aggregate? Is this not a modern version of Keynes’s paradox of thrift? In the context of humanity’s entire history could the recent invention of retirement be seen as an aberration from the norm (of working till you drop, or are supported by family)?

Tim Hodgson

Towers Watson

875 3rd Avenue, #17

New York, NY 10022

tim.hodgson@towerswatson.com

Extreme Risks and the Retirement Anomaly

Tim Hodgson

Retirement for the masses represents the briefest of anomalies over the broad sweep of human history. While the rich have always enjoyed their leisure, most of our forefathers and mothers worked until they were physically or mentally incapable of carrying on. The idea of self-sufficient time over many years in old age—what today we call retirement—is a very modern invention. To explore whether the concept of retirement is sustainable in the future, this chapter asks 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 modeling 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 likelihoods of extreme events. In this second case, it is less possible to diversify risk across time, hence one must weight more heavily the consequences 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. In particular, we focus on political, environmental, social, and technological risks, as these receive far less attention than the much more evident financial and economic risks. We conclude that extreme risks matter and deserve far more attention than given thus far. Partly as a consequence but also due to economic risks, retirement for the masses is also at risk.

Alternative World Views

Early developers of financial economics relied on economic theory, which itself drew on the mathematics and statistics of physics (Lo and Mueller 2010). Its initial inspiration was Newtonian physics with its cause-and-effect approach and linear relationships. In many economic models, therefore, the economy is assumed to be in, or moving toward, long-term equilibrium with stable parameters through time. Dynamics are permitted, with short-term cycles typically generated by external shocks to the system. Moreover, economics frequently uses the phrase '*ceteris paribus*', or 'all else being equal', but this cannot hold in a dynamic system. This mind-set has made it difficult to focus on rare 'tail events' and only recently are analysts devoting much attention to develop a theory of systemic risk.

By contrast, complexity science, or the study of complex adaptive systems,¹ has as its focus how system-level conduct emerges from interactions between components (c.f. Bostrom 2013). In particular, this approach explores how complex collective behavior can emerge if individual actors operate with simple rules and lacking central control. Moreover, signaling and information processing becomes key: a form of network arises when the actors produce and use information from internal and external sources. And finally, complex systems adapt and change their behavior over time, implying an absence of stable equilibrium.

Viewing financial markets as complex adaptive systems rather than as linear systems tending toward equilibrium has important implications (see Table 11-1). 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, efficient markets. Second, one must grapple with the 'interconnectedness of all things'.² For example, the collapse of Lehman Brothers in September 2008 shows that counterparty risk is subject to

nonlinear 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, modeling becomes much more complex. That is, shocks can be modeled with fatter tails and higher probabilities, but complexity science has not yet developed sufficiently to provide readily-usable quantitative models for finance and economics.

Table 11-1 here

The Irreversibility of Time

The conventional approach to think about risk considers all possible outcomes and weights them in accordance with their probability. In effect, the exercise freezes time, taking multiple copies of the world and running them forward as ‘parallel universes’. For instance, consider the following fair gamble: we will roll a fair die and stipulate that if I roll any number from one to five, I must pay you 50 percent of your current wealth including the present value of your future earnings.³ 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 die. The downside, paltry in comparison, is that if you roll a six you must pay me your entire wealth. The expected value of the bet is a 25 percent expected return.

While this seems attractive from an odds perspective, many would not take the bet, perhaps because people might think about it as follows. Instead of rolling the die once in each of six parallel universes, we stay in our familiar universe and roll the same die six times in

succession: the time average takes each of the six possible independent outcomes and makes them occur one after the other. Compound returns over the six periods and taking the sixth-root to calculate the per-period expected (time average) return: the time average is negative (−100 percent). In other words, the ensemble average *is* misleading; the 25 percent expected return unhelpfully disguises the meaningful (16.7 percent) likelihood that you lose everything. And when you have lost everything, you cannot go back and try again.

A more subtle point is that the expected return calculation effectively underweights the impact of extreme risk. To illustrate this, consider a world in which one invests a portfolio of financial assets with either a good or bad outcome. In the good case, one earns a return of 5 percent and this occurs most of the time; in the bad case, which only occurs once in the distribution, it causes severe or total loss. Table 11-2 shows that for, a number of runs of the thought experiment, the ensemble average return exceeds the time average return. The column pairs represent worlds with different probabilities for the single extreme outcome, starting with 1-in-1000 and moving right to 1-in-100. The rows show cases where the severity of portfolio loss rises from 99 percent in the first row to 100 percent in the bottom row. For all 20 combinations of probability and severity, the ensemble average return is always higher than the time average return, and in some cases significantly higher. Once the probability of the extreme event rises to 1-in-100 or higher, then 99 good outcomes of 5 percent are wiped out by a single extreme event. Finally, note the difference in the time average return between extreme losses of 99.999 percent and 100 percent: as in the die throwing experiment, once you have lost all your wealth, the game is over and your return is −100 percent, whether that occurs in the first or the last period. Losing 99.999 percent of your portfolio would clearly be painful, but the little that is left can then start

to grow again. This highlights the difference between an existential risk and a risk where ‘life’ continues to the next period, albeit in very poor shape.

Table 11-2 here

Of course one might object that a 99 percent or higher portfolio loss is too extreme to contemplate realistically. Yet some of the extreme risks discussed below might cause financial portfolios to become worthless, and history has offered several examples of entire stock markets losses. In the pension case, if the retirement fund were a defined benefit arrangement and the plan sponsor has ceased to exist, a large portfolio loss could mean that the fund’s mission failed: assets will run out before the liabilities are paid and, absent an insurance arrangement, some beneficiaries will receive nothing. So for them, at least, this would equate to a total portfolio loss. If the retirement fund were a defined contribution arrangement, a large portfolio loss would not qualify as existential since there is no contractual benefit to be broken. Instead, the adjustment is borne by individual participants, perhaps by accepting a lower retirement standard of living than anticipated. Even here, however, not all members are equal. A 50 percent loss for a 29-year old is fundamentally different to a 50 percent loss for a 59-year old; the older participant could perceive the loss as bordering on the existential. The practical takeaway is that avoiding, or reducing the probability of 100 percent (existential) losses is incredibly valuable.

Using a different lens to regard the world can increase our perception of the qualitative chances of an extreme event occurring,⁴ and embracing the reality of irreversible time meaningfully increases the significance of extreme risks. Next we consider what events might be considered extreme, in the context of their impact on retirement systems.

Identifying Extreme Risks

A large variety of risks could be included on a list of extreme events. Here we propose a framework or set of filters that can help identify which risks are of most import, and which can be safely ignored. Here we consider only first-order impacts and discuss them in isolation. We do, however, make an attempt 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. Moreover, context-specific risk assessment considers 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.⁵

Since our focus is on sustainability of the retirement system, we focus mainly on risks that can erode assets or increase liabilities; we also omit discussion of legal, process, and operational risks. Having acknowledged these limitations, our overall framework consists of six categories: finance, economics, politics, environment, society, and technology. We also touch only lightly financial and economic categories since these are widely commented on in other chapters in this volume (and Towers Watson 2010, 2011). Financial extreme risks would be sovereign default, banking crises, and insurance crises. Economic extreme risks would include depression, hyperinflation, currency crises, and the end of fiat money.

The six categories in our framework are not independent; in fact, some of the most interesting insights can be found at points of confluence; see Table 11-3 (and the Appendix). To be clear, some of these risks may seem quite extreme (and perhaps amusingly or ridiculously so). But our goal is to allow for ‘black swan’ or highly unusual events. Within a category, the risks

are listed in alphabetical order and do not represent any form of ranking. The risks listed are also intentionally extreme: in effect we ask ‘is it plausible?’ or, perhaps more accurately, ‘does our current state of knowledge suggest it is completely implausible?’ Next we evaluate, assuming this event occurred, what might be the consequences.

Table 11-3 here

Political. This category of extreme risks derives from policy decisions. In two cases, the link is direct: global trade collapse follows policy decisions to favor protectionism over openness and globalization, and World War III would flow from active decisions to declare war. For anarchy and political extremism, the link is less direct, yet poor prior policy decisions are likely to be necessary, if not sufficient, conditions for these risks to foment. Terrorism is included in the political category due to its ideological foundation, and because the target chosen for the act of terrorism is likely to have political ramifications.⁶

Environmental. These risks represent threats to human safety and well-being arising from a disruption to the planet’s environment. Two of these, alien invasion and cosmic threats, can be considered exogenous. While some might think that an alien invasion is too extreme to consider seriously, and its probability and consequences are unknowable, good risk management requires taking action in advance to protect against possible future consequences. Accordingly, the exercise is to scan the horizon with an open mind, and one can always apply additional filters later, to protect finite risk management resources. For now, all we can say is that too little is known to rule out the possibility of an alien invasion.

Two other environmental risks, biodiversity collapse and global temperature change, could be the result of human activity and thus they represent serious challenges given the

absence of effective global governance. Other possible sources of both risks are shown in Table 11-4.

Table 11-4 here

The final risk in this category is natural catastrophe. As earthquakes, for example, occur every day, the extreme version of this risk is either a confluence of extreme natural catastrophes (e.g. a magnitude 10 earthquake combined with a 25 meter tsunami, helped along by a Category Five windstorm), or the eruption of a supervolcano.

Social. Social extreme risks are threats that could adversely affect the smooth functioning of society. As noted above, the table categories are not independent and social risks link to policy decisions, the environment, and, in some cases, to technology. This is clearest in the case of a food/water/energy crisis which would have political, environmental, and technological drivers—as well as offsets.

Three of the risks are health-related. Pandemics are a favorite of commentators on extreme risks, as there are good data (at least in relative terms). Here we postulate that a new disease agent hits the ‘disease sweet spot’ of both high infectivity and high mortality (in practice these have often been trade-offs). Health progress backfire refers to a reversal in the trend of improved health, while working in the other direction, extreme longevity becomes a risk when viewed through the lens of a retirement provider. The final risk identified in this category is organized crime, to the extent that legitimate economic activity ceases to be viable in a major country or region.

Technological. Our last category of extreme risks concerns technology. Here one might contemplate a failure in current technology (nuclear contamination or infrastructure failure); possible consequences of emerging technology (cyber warfare and biotech catastrophe); or a

technological singularity—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.

Assessing Extreme Risks

Next we seek to determine which of these risks might be more or less material in the context of a retirement system. The traditional risk assessment approach would be to learn from historical data; a qualitative overlay can be added, if for example, a structural change is believed to have occurred. But when thinking about extreme risks, by definition there may be no or very small historical evidence.

For this reason, we take a deliberately qualitative (deductive) approach to the assessment process. In the first phase, a team of Towers Watson researchers reviewed the research literature and historical data on past extreme events. Team members then independently generated scores. In phase two, the independent scores were compared and debated, with a single consolidated scoring approach generated. For stage three, the consolidated scores were reviewed by a senior committee and further refinements suggested. Hence what we propose has been subjected to a rigorous oversight process, making our conclusions robust albeit qualitative.

Next we score each of the extreme risks in Table 11-4, and results are illustrated graphically in Figure 11-1. 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 draw on and adapt the qualitative risk categories of Bostrom (2013). These scores relate to the most extremely negative manifestation of the risk. For instance, the consequences of earth being visited by aliens span a massive range from the beneficial to the extinction of the human species.

As we are concerned here with extreme negative risks, it is the latter potential outcome we score. Also, rather than assign each risk a probability (the quantitative route), we assign it one of four categories representing a likelihood of occurrence of 1-in-10 years, 1-in-20 years, 1-in-100 years, and less than 1-in-100 years. (For the technically-minded, this is akin to a high-alpha power law for the distribution and it implies we believe these events should be considered far more likely than if a normal distribution were used).

Figure 11-1 here

We split the potential impact of each risk into two separate dimensions, namely the intensity, and the scope (or geographical/temporal spread). The intensity is assigned to one of three states labeled ‘endurable’, ‘crushing’, and ‘existential’.⁷ 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 affecting more than one generation, but that would then fade or reverse. Pan-generational refers to an impact that would affect all subsequent or all previously potential generations (such as extinction of the human species). Of course there is a danger that these two dimensions are not independent; for instance, a food/water/energy crisis could 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, assessed as low, medium, or high. In the graphical representation (Figure 11-1), 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.⁸ Figure 11-1 is perhaps more useful for assessing the overall findings. One might have hoped that the most likely risks were endurable, while the existential risks less likely, and the Figure does show this general shape. One might have also anticipated that local impact was more frequent than global, global more frequent than trans-generational, and pan-generation the least frequent; this too is the case in our view.

Of course in the retirement context, we are most interested in impacts on assets and liabilities. If we split Figure 11-1 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 as assets are likely to be worthless and the value of the liabilities may well have fallen to zero too. The second group comprises the seven points to the left with crushing intensity of impact. For these events, the effect on assets will be global and materially negative. There will 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. That is, if your nation is at war, you lose lives and your liabilities go down. If your neighbors are neutral, their liabilities do 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 seven points (endurable intensity). These are less homogenous, but in general, we would expect a more muted impact on assets and liabilities. The clear outlier in this group is extreme longevity which would explicitly increase the liabilities.

Association. As noted above, these risks are not entirely independent, so we have also developed an ‘association’ matrix (see Table 11-5). This is not a correlation matrix: correlations require data to calculate and even then they say nothing about causality. Instead, here 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 evaluating 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. ‘L’ for low means that we believe X could cause Y, or is a contributory factor. ‘H’ for high means we believe the causality is material, so X is likely to, or will, cause Y. For example, reading across the first row of entries shows that we believe anarchy (P1) *could* cause, or contribute to, global trade collapse (P2), World War III (P5), a food/water/energy crisis (S2), biotech catastrophe (T1), infrastructure failure (T3), and nuclear contamination (T4). Anarchy *is likely to* cause political extremism (P3), and a sharp rise in organized crime (S4).

Table 11-5 here

The table can also be read down the columns. Here 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). We emphasize that this is an association matrix between only the risks we are considering, and therefore it is silent on whether other contributory factors outside our current consideration also matter.

Many comments could be made about the associations, but here we focus on the key ones for our purpose. The environmental risks (rows E1 to E5) are powerfully causal, so contingent on one of them occurring, other extreme risks could be triggered. By contrast, the social extreme risks (rows S1 to S5) have low causal power: they are unlikely to trigger further extreme events. Yet the social risks could have multiple triggers, as indicated by the number of entries in columns S2 to S5. This is especially the case for food/water/energy crisis (S2), attributable to our

assessment that the demand for water and food may soon start to exceed supply. With little current buffer, it would take little to tip the balance the wrong way. We also note that there are some truly independent risks in the table. Columns E1, E3, S1, and T5 imply that an alien invasion, cosmic threats, extreme longevity, and technological singularity are unlikely to be triggered by any of the other extreme events considered here.

Next we create a ranking of the importance of the risks; this introduces no new information, but rather combines the scores for each risk into a single number that can then be ranked. The approach is straightforward: the more likely a risk, the higher the ranking it receives. Likewise, the more certain a risk, the greater its intensity of impact; the larger the scope of the risk impact, the higher its ranking. The scoring necessarily involves some element of rough justice, in that it requires trading off between, say, likelihood and intensity. But instead of focusing overly much on precise ranking, it is more useful to arrive at an overall judgment as to whether the ranking is satisfactory, inquiring whether the risks at the top and at the bottom seem reasonable. Results appear in Table 11-6.

Table 11-6 here

At the top of the list is a food/water/energy crisis (S2), driven by our assessment that this is one of the most likely risks; hence there is *relatively* little uncertainty attached to either its likelihood or the consequences. The consequences themselves, locally crushing, are not globally severe (in relative terms), so these do not imply a top ranking. In contrast, the second ranked risk, global temperature change (E4), has much more severe consequences (trans-generationally crushing) but it 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 in the bottom 10 overall, with alien

invasion (E1) ranked at 19. This is because their poor consequences are diluted by our assessment of low likelihood and relatively high uncertainty. At the very bottom is political extremism (P3), where the locally-crushing consequences are not particularly severe, the likelihood is relatively low, but there is a high degree of uncertainty of how assets and liabilities would be affected.

As noted above, different scoring systems and weights could produce somewhat different rankings, but the power of our approach is that it combines and trades off the four risk scores in a consistent manner. Overall, the value of this approach is to challenge pre-conceptions and highlight which risks to prioritize, when it comes to management actions.

Implications for Retirement and Retirement Systems

Above we showed that adequate saving is necessary but not sufficient for a comfortable retirement period. And extreme risks pose a threat to the quality of retirement in terms of security and/or wellbeing, whether political, environmental, sociological, or technological. Similarly, extreme risks pose a risk to retirement systems. To protect against these, we require more robust investment portfolios and balance sheets. Accordingly, we consider how to incorporate the framework developed above in the process of managing a retirement fund's balance sheet.

One option would be to slightly reduce expected returns, or push up volatilities and/or correlations, so as to incorporate the impact of infrequent extreme events (Towers Watson 2011). Another approach is to model dynamic switching: here one must make explicit two sets of assumptions ('normal' and 'extreme'), or one could design a second, extreme-risk, portfolio directly (Moore and Pedersen 2014). A third option is to construct a hedging overlay, though

hedging comes with its own set of problems including high carrying cost and often the use of derivatives. Not all extreme risks can be hedged, and hedges are imprecise. In fact, once one moves beyond financial and economic extreme risks to confront the extreme risks described above, hedging becomes very difficult indeed. Whether the counterparties would be willing to and capable of paying out if the extreme event happened is also unclear.

Retirement systems have three broad hedging strategies. The first is to hold cash; this is likely the ‘oldest, easiest, and most underrated source of tail risk protection’ (Montier 2011) and is closest to a generic hedge for all of the risks considered in this chapter. Over long historical periods, in many countries, cash has held its real value through both episodes of deflation and inflation but there is no guarantee that this will be the case in the future. Cash is likely to be effective in partially hedging risk up to and including crushing intensity, but it is likely to become worthless once we cross over to existential intensity (all other assets are also likely to be worthless, but the funding ratio could be quite good if all the liabilities have been extinguished). Also there could be non-linear breakpoints between crushing intensity and existential intensity. Another consideration regarding holding cash is which currency (or currencies) to favor. Since retirement system liabilities tend to be domestic, there is a natural draw to domestic currency cash. Yet this represents a relatively concentrated risk position and so a degree of diversification may be preferable.

A second hedging strategy for retirement systems is to use derivatives, though cost and usefulness are often in opposition. The cost of derivatives protection can often be reduced by specifying more precise conditions for the payout, but the more precise the conditions, the greater the chance that these will not be met exactly so the protection will not pay. Derivatives in the form of swaps are an obvious way to hedge the risk of extreme longevity, but capacity is

likely to be severely constrained relative to potential demand (c.f. Coughlan 2014). Beyond longevity hedging, derivatives could possibly be used to help with one or two of the risks such as global trade collapse, or commodities futures could be used to hedge food/water/energy crisis. In the main, however, derivatives are probably better suited to hedging financial and economic risks.

A third retirement system hedging strategy is to hold assets that are negatively correlated with their liabilities. Typical hedging assets include the highest quality sovereign nominal and inflation-linked bonds, as well as precious metals. More creativity is probably required to handle some of the risks. For example, with respect to World War III, one strategy would be to hold assets in neutral countries; of course this would require foresight and willingness to hold a portfolio very different from that of peers. In fact, all of the political risks require some sort of portfolio-exclusion approach, avoiding assets in jurisdictions likely to be affected. In any event, no single asset is likely to protect against all possible bad outcomes. Moreover, there is no guarantee that the expected performance of the hedge asset will actually transpire in the event.

To make retirement funds more resilient in the face of extreme risks, we suggest they undertake a prioritization exercise. At the top of the list is to focus on the events that might permanently impair the investor's mission, to identify which extreme risks matter and which can be ignored. For the former, insurance is indicated since the investor cannot afford to self-insure. For the latter, an investor should do the 'simple' things, including diversifying the portfolio across as many return drivers as possible; today many institutional portfolios are heavily concentrated in equity or growth risk. Additionally it is useful to diversify within asset classes; today most pension funds are heavily exposed to domestic sovereign bonds for liability-matching purposes. Extreme risk thinking would suggest reducing the quality of the match to reduce the

risk concentration. Holding cash strategically also provides optionality. And finally, over time, greater complexity can be added 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 inexpensive 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 percent of earnings, and earn a real return of 3.5 percent 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 percent 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 percent of GDP (Towers Watson 2013), implying much too few resources available to support retirement. So we have either chronically undersaved for retirement and need to do better, or other assets outside the pension system may be the answer. If the diagnosis is chronic undersaving, 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 chapter 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 political, environmental, social, and technological 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 risks are difficult to hedge, a relatively high cash weighting appears versatile and effective. Perhaps rather than changing investment strategy, we must save more and build a larger risk buffer, a course of action currently forced onto banks.

Issues we lacked space to discuss include two worth mentioning. First, embedded in our likelihood scores is the fact that the world lacks effective global governance, ‘G0’.⁹ Many political, social, technological, and two of the environmental risks identified are clearly exacerbated by an absence of effective global governance. Conversely, such governance could materially reduce the likelihood of these multiple risks. Second, we noted the possibility that technology could ‘run amok.’ Naturally, technology has been of substantial benefit to humankind, but many unintended consequences arise 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.

Acknowledgements

The author acknowledges his colleague, Liang Yin, for support in this research, though any errors and omissions are those of the author alone.

References

- Adams, D. (1987). *Dirk Gently's Holistic Detective Agency*. London, U.K.: William Heinemann Ltd.
- Adams, D. (1991). *The Long Dark Tea-Time of the Soul*. New York, NY: 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, U.K.: Oxford University Press, pp. 504-519.
- Coughlan, G. (2013). 'Longevity Risk Management and Corporate Finance,' PRC Working Paper 2013-20. Philadelphia, PA: Pension Research Council.
- Göbel, C., and L. H. Ong (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, eds., *Global Catastrophic Risks*. Oxford, U.K.: 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, No. 29, 3rd ed. Ithaca, NY: Cornell University.

- 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*. New York, NY: Basic Books.
- Montier, J. (2011). 'A Value Investor's Perspective on Tail Risk Protection: An Ode to the Joy of Cash,' GMO white paper.
- Moore, J., and N. K. Pederson (2013). 'Implications for Long Term Investors of the Shifting Distribution of Capital Market Returns,' PRC Working Paper 2013-24. Philadelphia, PA: Pension Research Council.
- National Aeronautics and Space Administration (NASA) (2013). *Kepler: A Search for Habitable Planets* [website].
<<http://kepler.nasa.gov/>>
- Petranek, S. (2002). *10 Ways the World Could End* [Video file].
<http://www.ted.com/talks/stephen_petranek_counts_down_to_armageddon.html>
- Rees, M. (2003). *Our Final Century: The 50/50 Threat to Humanity's Survival: Will the Human Race Survive the Twenty-first Century?* London, U.K.: William Heinemann Ltd.
- Santa Fe Institute. *History of the Santa Fe Institute* [website]. Santa Fe: Santa Fe Institute.
<<http://www.santafe.edu/about/history/>>
- Soros, G. (2009). 'One Way to Stop Bear Raids,' *The Wall Street Journal*. March 24.
- 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. London, U.K.: Towers Watson.

Towers Watson (2011). *Extreme Risks – The 2011 Update*, TW-2011-INV-00130. London, U.K.: Towers Watson.

Towers Watson (2012). *The Impossibility of Pensions*, TW-EU-2012-25855. London, U.K.: Towers Watson.

Towers Watson (2013). *Global Pension Asset Study 2013*. London, U.K.: Towers Watson.

United Kingdom Home Office (2013). *Reducing Knife, Gun and Gang Crime* [electronic publication]. London, U.K.: Home Office.

<<https://www.gov.uk/government/policies/reducing-knife-gun-and-gang-crime>>

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,’ [lecture]. Santa Fe, AZ: Santa Fe Institute.

Wikipedia (2013a). ‘September 11 Attacks,’ [website]. Accessed April 2013.

<en.wikipedia.org/wiki/September_11_attacks>

Wikipedia (2013b). ‘Biodiversity,’ [website]. Accessed April 2013.

<<http://en.wikipedia.org/wiki/Biodiversity>>

Wikipedia (2013c). ‘Nuclear and Radiation Accidents,’ [website]. Accessed April 2013.

<http://en.wikipedia.org/wiki/Nuclear_and_radiation_accidents>

Wikipedia (2013d). ‘Grey Goo,’ [website]. Accessed April 2013.

<http://en.wikipedia.org/wiki/Grey_goo>

World Economic Forum (WEF) (2013). *Global Risks 2013 – An Initiative of the Risk Response Network*. Geneva, Switzerland: WEF.

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, American Geophysical Union. 23(8): 837–840.

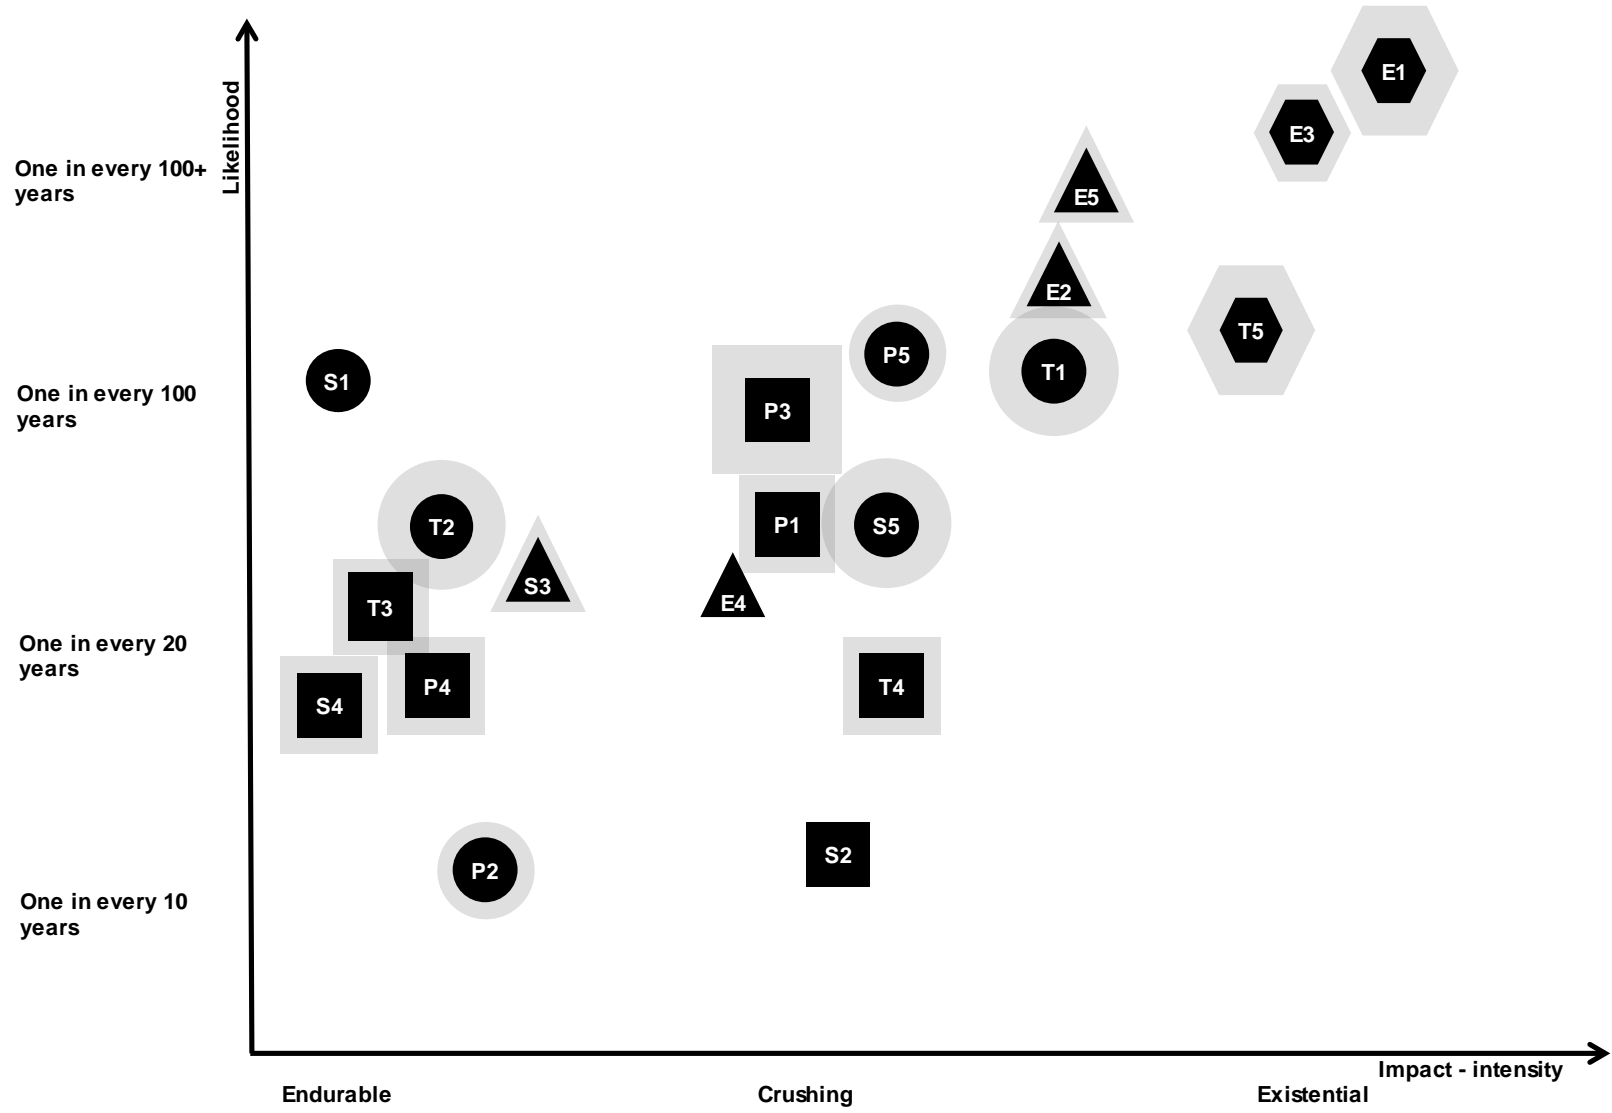


Figure 11-1. Extreme risks scoring. *Notes:* the scope of the impact is shown in different shapes (Local – Square; Global – Circle; Trans-generational – Triangle; Pan-generational – Hexagon). The size of the shape is indicative of the level of uncertainty with regards to the likelihood/impact of each risk.

P1 Anarchy	E1 Alien invasion	S1 Extreme longevity	T1 Biotech catastrophe
P2 Global trade collapse	E2 Biodiversity collapse	S2 Food/water/energy crisis	T2 Cyber warfare
P3 Political extremism	E3 Cosmic threats	S3 Health progress backfire	T3 Infrastructure failure
P4 Terrorism	E4 Global temperature change	S4 Organised crime	T4 Nuclear contamination
P5 World War III	E5 Natural catastrophe	S5 Pandemic	T5 Technological singularity

Source: Author's analysis.

Table 11-1 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.
Coupled / interacting	Transactions cause a security's price to change, which triggers subsequent transactions and further price changes. In addition, news flow can cause coupling when unrelated parties trade on the same story.
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 and Hudson 2004).
Unintended	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-

consequences	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 returns were i.i.d. (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 2009).
Multiple phases (regimes)	Academics have identified different market regimes in monthly returns data (Kritzman and Li 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 CFO of Goldman Sachs, explained that ‘we were seeing things that were 25-standard deviation moves, several days in a row’ (Larsen 2007). 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.

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 by dint of their reflexive nature appear more likely to be complex than complicated

Notes:

^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 (West 2011).

^b This column indicates how financial markets exhibit the characteristics of complex adaptive systems.

Source: West (2011) and the author.

Table 11-2 Expected Returns (%) 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	-00	4.58	-100	3.95	-100

Notes:

The ensemble average is calculated using the formula: $E(r) = \sum_{i=1}^N \frac{r(i)}{N}$

The time average is calculated using the formula: $T(r) = \prod_{i=1}^N [1 + r(i)]^{1/N} - 1$

Source: Author's calculations.

Table 11-3 Extreme Risks by Category

Political	
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 (e.g. 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).
Environmental	
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 $>1,000 \text{ km}^3$ 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	Major shortfall in the supply of food, water, or energy causing political strife and widespread death and severe

crisis	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 (e.g. 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 human behaviors, 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's analysis.

Table 11-4 Extreme Risk Scoring

	Likelihood 1 – one in 10 years 2 – one in 20 years 3 – one in 100 years 4 – one in 100+ years	Uncertainty A degree of High (H), Medium (M), or Low (L)	Impact – intensity 1 – Endurable 2 – Crushing 3 – Existential	Impact – scope 1 – Local 2 – Global 3 – Trans-generational 4 – Pan-generational
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
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

Source: Author's analysis.

Table 11-5 Extreme Risk Association Matrix

	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	
P1 Anarchy		L	H		L							L		H		L		L	L		P1
P2 Global trade collapse	H		L		L							H		H			L	L			P2
P3 Political extremism		L		L	H							L					L		L		P3
P4 Terrorism	L		L													L		H	L		P4
P5 World War III		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				H					H	L		L			L			E4

E5 Natural catastrophe	H	H					H		H			H						H			E5
S1 Extreme longevity												L									S1
S2 Food/water/energy crisis	H				L									L	L						S2
S3 Health progress backfire																L					S3
S4 Organised crime	L																				S4
S5 Pandemic		L																			S5
T1 Biotech catastrophe	H	H										L	L		L						T1
T2 Cyber warfare					L													H	L		T2
T3 Infrastructure failure												L									T3
T4 Nuclear contamination												L						L			T3
T5 Technological singularity	L	L	L																		T5
	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’s analysis.

Table 11-6 Extreme Risk Ranking

1.	S2 Food/water/energy crisis
2.	E4 Global temperature change
3.	P2 Global trade collapse
4.	S3 Health progress backfire
5.	T4 Nuclear contamination
6.	S1 Extreme longevity
7.	P4 Terrorism
8.	T3 Infrastructure failure
9.	E2 Biodiversity collapse
10.	S5 Pandemic
11.	P5 World War III
12.	T5 Technological singularity
13.	E3 Cosmic threats
14.	T2 Cyber warfare
15.	P1 Anarchy
16.	S4 Organised crime
17.	E5 Natural catastrophe
18.	T1 Biotech catastrophe
19.	E1 Alien invasion
20.	P3 Political extremism

Source: Author's analysis.

Appendix: A Taxonomy of Extreme Risks

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

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 (Göbel 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 percent (Caplan 2006).

P4 Terrorism. 9/11 caused almost 3,000 deaths and the Dow Jones Industrial Average index to fall by more than 14 percent within the first week of market reopening. New York City's GDP was estimated to have lost \$27.3 billion for 2001-2012 (Wikipedia 2013a). 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) (NASA 2013), 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 percent 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 (Wikipedia 2013*b*). 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 percent 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 (e.g., 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 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 (e.g., 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 banes 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 freshwater 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 warns (WEF 2013) 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 (2002). 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 U.K. Home Office (2013) suggests that organized crime costs the U.K. between £20 and £40 billion each year, 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 form 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 (e.g., 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 casualties. For example, it is believed that Plague of Justinian in AD 541-542 killed 50 percent 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 1918-1919 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 (e.g., 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 (e.g., cyber-attack), natural disasters (e.g., earthquake or flooding), or even cosmic threats (e.g., 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 \$7 billion of property (Wikipedia 2013c). 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 (Wikipedia 2013*d*). 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.

Endnotes

¹ See for instance the Santa Fe Institute, <http://www.santafe.edu/about/history/> (Santa Fe Institute 2013).

² A phrase borrowed from Dirk Gently, a fictional ‘holistic’ detective who believed all things were connected in some way (Adams 1987, 1991).

³ This is a thought experiment so we will gloss over my ability to pay; assume my credit is pristine.

⁴ 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.

⁵ Clustering of risks is the basis of scenario analysis which should be done; this is outside the scope of the present chapter.

⁶ We acknowledge that terrorism is a weekly, if not daily, occurrence somewhere around the world; the extreme risk here would be a terrorist act comparable to, or worse than, 9/11.

⁷ We do not define these states, as the labels are self-explanatory. But one might think about oneself in the three states as follows: an endurable risk could represent a broken leg; a crushing risk might imply the loss of a limb or paralysis; and an existential risk might refer to the loss of self-awareness or loss of life.

⁸ An extra layer of sophistication could have added by extending the shapes in either the vertical or horizontal direction, implying greater certainty in one dimension and less in the other. But there is no reason to stop there; one could also move the border off-center to show that the uncertainty is asymmetric. In view of the high intrinsic uncertainty involved when considering extreme risks, this extra sophistication would likely provide the impression of spurious accuracy and give a false impression of the level of signal relative to the noise.

⁹ A term and concept coined by Ian Bremmer of Eurasia Group.