Prepared by:

Neil Cantle, Milliman Limited
MA, ASA, FIA

Neil Allan, University of Bristol
Fellow of Bristol Systems Centre

August 2009

Milliman

Insurance Industry Under the Microscope

Complexity-based Analysis Reveals What Is Beneath the Surface


Milliman

TABLE OF CONTENTS

ACKNOWLEDGEMENTS

The complexity analysis in this report was carried out using DACORD™, a tool developed in collaboration with DRTS Ltd.® (www.drts.co.uk).

EXECUTIVE SUMMARY

Like most service businesses, the insurance sector is built on perceptions and a social construct of shared risk. The concept of insurance has emerged from the desire for individuals to take risks they prefer not to bear alone. Insurance provides a mechanism through which such risks can be spread across society. We can think of a sector, a company, or an entire economy as a *system*, which contains rules, cultural norms, principles, and processes.

Traditional approaches to analysing performance generally look at different aspects of the performance and measure them individually. There is then a process of recombining them, using ratios or metrics to form an overall assessment of whether things are going well or not. But this approach misses the complex interactions between factors; in addition, the importance of some of the factors can change and adapt over time. Indeed, much of the underlying behaviour arises from the combinations of interactions, so there is little to learn from studying the individual components.

By the conceptualisation of risks using complex systems theory, Milliman has developed an approach that enables risk managers to look below the surface to see the indications of trouble before it breaks. Milliman's CRisALIS™ methodology and techniques support risk managers in making early decisions about emerging risks, mitigating the worst effects.

Milliman's CRisALIS™ methodology and techniques support risk managers in making early decisions about emerging risks, mitigating the worst effects.

One critical metric we use to assess the state of the industry, and the companies within it, is the level of *uncertainty*. This metric should be running at a level for companies to be sufficiently complex that they can generate a good level of performance, but it should not be so high that they can no longer manage the business.

The financial services sector is primarily based around packages of services so there are a significant number of interactions required to deliver the output. Given the large number of participants in delivering the final product, and hence the large number of interfaces and opportunities for *variation* in performance to occur, it therefore makes sense that the industry, and the companies within it, should normally be operating at a relatively high level of *uncertainty*.

In this paper, we look at three levels of analysis to assess the performance of the life assurance sector:

- The behaviours of the life assurance sector as a whole
- Whether the size of an organisation matters in looking at its vulnerability to performance issues
- Whether there is difference in the performance of proprietary companies and those owned by members (mutuals, friendly societies, cooperatives, etc.)

Our analysis of the UK life insurance sector has shown that the headline performance indicators on which many analysts rely do not always tell the whole story. In the end they may show poor results, but often take too much time to integrate *bad news*. By using a combination of complexity-based tools, it is possible to look deeper into the information about the sector's performance and seek out the patterns that truly signify its current state.

By using a combination of complexity-based tools, it is possible to look deeper into the information about the sector's performance and seek out the patterns that truly signify its current state.

INTRODUCTION

It's All About People

The insurance sector, like most service businesses, is built on perceptions and a social construct of shared risk. It has emerged through the desire for individuals to take risks which they prefer not to bear alone. The sector provides a mechanism through which such risks can be spread throughout society. We can think of a sector, company, or an economy as a *system* that contains rules, cultural norms, principles, processes, and so on. A key component of this system is people—they are the *moving parts* in this system, they are an adaptive feature that can lead to extreme non-linear behaviour, such as groupthink, greed, and fear, each reinforcing the next.

The insurance sector organises itself into company structures that set goals and lay down processes and principles for staff to follow. Companies also generate cultural norms through the actions and beliefs of senior management. People choose to join companies and then go about their allocated tasks, framed in the context of these constraints.

Systems involving people are capable of generating complex organisations with complex products and services. Even though individuals follow a relatively simple set of rules, their interactions and subsequent adaptation and evolution result in an overall system behaviour that is far from simple.

A key feature of such systems is that understanding the behaviour of the individual components tells you little about the behaviour of the system as a whole. Failing to capture the interactions between the sub-components is where traditional approaches to risk and performance measurement fall down. It is the complex, adaptive set of interactions that give rise to the organisation's performance, the sector's performance, and also the emerging threats. Moreover, the complex interactions are primarily initiated by people.

It is the complex, adaptive set of interactions that give rise to the organisation's performance, the sector's performance, and also the emerging threats.

Signs of Complexity

Traditional approaches to analysing performance generally look at different aspects of the performance and measure them individually. There is then a process of recombining them, using ratios or metrics to form an overall assessment of whether things are going well or not.

This approach misses the complex interactions between factors and that the importance of some of the factors can change and adapt over time. Indeed, the underlying behaviour arises from the combinations of interactions, so there is little to learn from studying the individual components.

Milliman's CRisALIS approach looks at the whole system, and lets the interactions indicate which measures are important to its performance and how they relate. This is a bit like a headache, running nose, and sore throat letting you know that you probably have a cold virus rather than three separate ailments.

From the study of complex systems theory, we know that, for a healthy system, there is a link between the level of complexity and performance capability. An increase in complexity allows for an increase of performance, up to a maximum threshold. This threshold is sometimes referred to as *the edge of chaos* and is the point at which the system essentially undergoes a change in mode. This change is often very sudden and its timing difficult to predict. Because real-world measurement is often only available with time lags, it is unwise to get too close to the threshold; the actual position may have moved across the boundary since the previous measurement. In practice we therefore know that a position *close* to the boundary is highly dangerous. We also know that over time the system adapts and the level of maximum complexity it can cope with will change. A key system measure therefore is the level of complexity relative to this threshold.

From the study of complex systems theory, we know that, for a healthy system, there is a link between the level of complexity and performance capability.

This study shows that the headline performance metrics can be misleading but that, if you know how to interpret the signs, they do contain information within their interconnectedness about the build-up of the underlying system performance.

Terminology

One indicator of complexity we look for is *uncertainty*. Simply speaking, the more uncertainty a system has, the less easy it is to control because there is high unpredictability about what it will do next. Uncertainty is a measure used to assess how easy it is to know what a system will do next based upon knowledge of its progress so far. According to information theory, the amount of information in an observation x is $-\log p(x)$ where p is the probability of x being the information we want. We then look at the average amount of information in the organization's performance variables and this is equivalent to looking at the uncertainty removed after seeing the actual performance of the organization. So when we are perfectly certain about what will happen next we learn nothing by watching the actual performance and our uncertainty is zero. When it is perfectly unclear to us what will happen next our uncertainty is 1.

Uncertainty is a measure used to assess how easy it is to know what a system will do next based upon knowledge of its progress so far.

By looking at sets of variables that describe the system, we can observe when and how information is being shared between those variables, which gives a measure of uncertainty or how easy it is to predict what might happen next. But when variables are no longer communicating, that is more evidence of unpredictable future behaviour. The metric we use essentially assesses how knowable the system's current state is based upon knowledge of the previous states.

There is a maximum amount of uncertainty with which each system can cope before it becomes unmanageable. This maximum changes over time and can be deduced by examining performance data of the system. We compute the *Relative Uncertainty* as being the current uncertainty relative to the maximum. As the measure approaches 100% (and in reality more like 90%–95%) we know that the system will become unstable and suffer a corrective event. In practice, this event corresponds to a dramatic reorganisation of the system, such as a major risk event or a merger or other significant restructuring. Note that this event is still actually a continuous process where the unfolding risk situation is fluid and hard to predict, but it is a point in time that most people would point to as the trigger occurrence of a risk. We can therefore use *Relative Uncertainty* as one indicator of the complexity of the system.

In order to understand more about where this build-up of uncertainty comes from we can also look at a broader measure of complexity that takes into account more factors influencing a system's complexity. Other factors that influence complexity are: the system structure; number of variables, levels of interdependencies, feedback loops, and the extent to which the system structure is exhibiting a power-law distribution. For our analysis we compute a complexity metric that takes account of these factors in addition to uncertainty where we have multiple data points describing the system. These additional factors give us information about the rigidity of the system. The more rigid the system becomes, the less easy it is to adapt to new circumstances and the more likely it is to become sensitive to changes in environment that require adaptation. Also, the more rigid the system is the more rapidly and uncontrollably risks can propagate through the system.

Analysing and understanding the drivers of this complexity metric therefore gives us insight into where the underlying source of emerging risk may be and gives us indications of the emerging pattern of events so we may know where to deploy risk management actions most effectively. When we find a system exhibiting high complexity, complexity science tells us that the system is in a highly vulnerable state and is highly susceptible to changes in its environment, making it prone to failure or loss.

Analysing and understanding the drivers of this complexity metric therefore gives us insight into where the underlying source of emerging risk may be and gives us indications of the emerging pattern of events so we may know where to deploy risk management actions most effectively.

OVERVIEW OF THE LIFE ASSURANCE SECTOR

We first look at the traditional headline performance metrics of the sector since 2000, and then view the same performance indicators through the lens of complexity measures.

Headline Performance

Since 2000 the insurance sector has suffered two market crashes. The first crash came in 2000 as the bull run from October 1998 finally ended. The second started in mid-2007 and has brought the markets back down to the levels they fell to in January 2003 during the final stages of the 2000 crash.

During the period 2000–2007, the life assurance sector has actually fared reasonably well. The first crash exposed weaknesses in some strategies and caused companies to focus on rebuilding and preserving capital. The introduction of more robust capital management practices has certainly left companies in better shape to weather the current storm. That is not to say, of course, that things are in good shape—simply that they are improved.

The chart in Figure 1 shows the relative performance of some key industry indicators for the sector as a whole.

FIGURE 1: RELATIVE PERFORMANCE OF KEY INDICATORS SINCE 2000

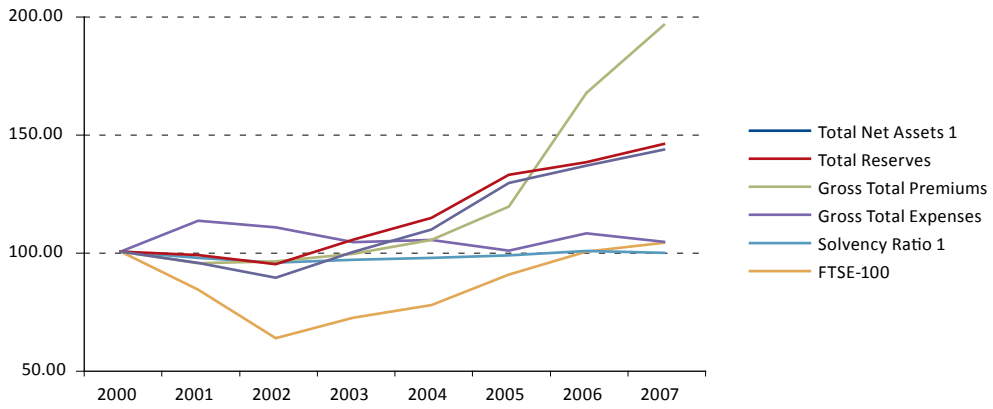

Source: FSA returns from SynThesys, FTSE-100

Looking at the headline figures, one might conclude that, since the 2000 crash, the industry has grown quite well, has kept expenses relatively under control, and has grown premiums faster than might be expected, while also returning its solvency ratio to a more healthy position.

However, looking at this performance from a different perspective we can see that the headline growth is not as robust as it may appear.

One of the key metrics we use to assess the state of the industry and the companies within it is the level of uncertainty.

A Different Perspective

One of the key metrics we use to assess the state of the industry and the companies within it is the level of uncertainty. As outlined in the introduction, the metric should be running at a level for companies to be sufficiently complex that they can generate a good level of performance, but it should not be so high that they can no longer manage the business. The financial services sector is primarily based around packages of services so there are a significant number of interactions required to deliver the output. Given the large number of participants involved in delivering the final product, and hence the large number of interfaces and opportunities for variation in performance to occur, it therefore makes sense that the industry, and the companies within it, should normally be operating at a relatively high level of uncertainty.

As can be seen in Figure 2, the industry seems to operate at a level of over 85% of the maximum level that it can manage. At this level the complexity of the industry is high enough to permit good performance but not so high that every participant in the industry is close to its maximum threshold and hence not precariously exposed to environmental shifts.

FIGURE 2 : TOTAL RELATIVE UNCERTAINTY OF LIFE SECTOR

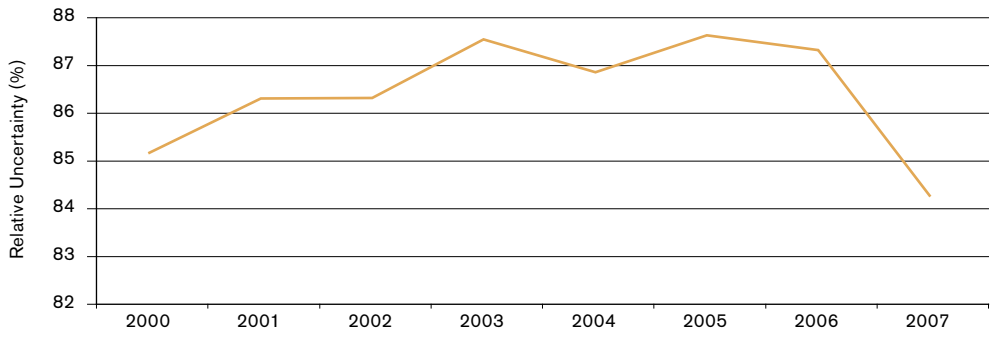

Source: Milliman analysis using DACORD™ (DRTS®)

When a system reaches a critical threshold it tends to become unstable and is usually associated with some kind of large risk emerging. It can be seen that, during the 2000–2003 turmoil, the level of uncertainty increased dramatically until it finally collapsed during 2003 as the industry system restructured itself. It is interesting to note that uncertainty continued to rise after 2004, which suggests that the apparently healthy recovery we saw in Figure 1 (on page 5) is less robust than it appears using traditional measures. The emergence of the current economic crisis is blatantly clear, with a dramatic release of uncertainty between 2006 and 2007 as it becomes evident from the information that the sector is in trouble and it enters a period of restructuring and risk mitigation.

The emergence of the current economic crisis is blatantly clear, with a dramatic release of uncertainty between 2006 and 2007 as it becomes evident from the information that the sector is in trouble and it enters a period of restructuring and risk mitigation.

Looking a little more closely at where this collapse in uncertainty comes from, we can investigate the relative levels of four key performance indicators, as shown in Figure 3.

FIGURE 3 : RELATIVE UNCERTAINTY OF KEY INDICATORS

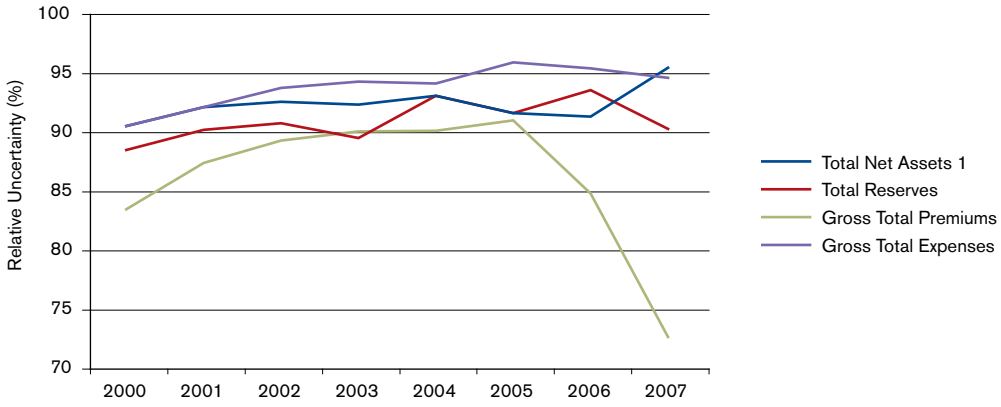

Source: Milliman analysis using DACORD (DRTS)

This analysis reveals that the dramatic drop in uncertainty is mainly caused by premiums. It is also interesting to note that there is a large increase in the uncertainty around total net assets. This picks up on the fact that the current economic turmoil has led to a large degree of volatility and difficulty in valuing assets, but could also be indicative of the build-up of the current economic situation before it broke.

The information contained in the FSA returns can be considered as multiple indicators telling us about different aspects of each company's performance. In aggregate they give us some information about the sector as a whole. As we have remarked, it is the interaction between components of organisations and industries that give rise to impacts on system performance, good and bad. We do this primarily by looking at the how the different indicators are sharing information with each other; this gives us a measure of their connectedness. People often refer to indicators that share a pattern of behaviour as being correlated, but clearly there is something much deeper occurring because these apparently correlated behaviours change over time.

The current economic turmoil has led to a large degree of volatility and difficulty in valuing assets, but could also be indicative of the build-up of the current economic situation before it broke.

The diagram in Figure 4 shows how a set of variables from the FSA returns are connected. The variables are shown along the diagonal, where blue nodes indicate variables that are particularly important to the performance of the sector at that time. Bolder lines indicate variables that are sharing more information with each other at that time. Information sharing is deeper than correlation, as it indicates that the variables are actually connected in some way at this point in time and are influenced by each other; correlation by contrast simply observes that they have similar behaviour. In computational terms, *information sharing* looks at the distance between the joint distribution of two variables and their product distribution. When two variables are independent the information shared is zero.

Figure 4 refers to the status in 2000 and therefore reflects a small impact of the crash that began that year.

FIGURE 4: INFORMATION SHARING BETWEEN KEY INDICATORS IN 2000

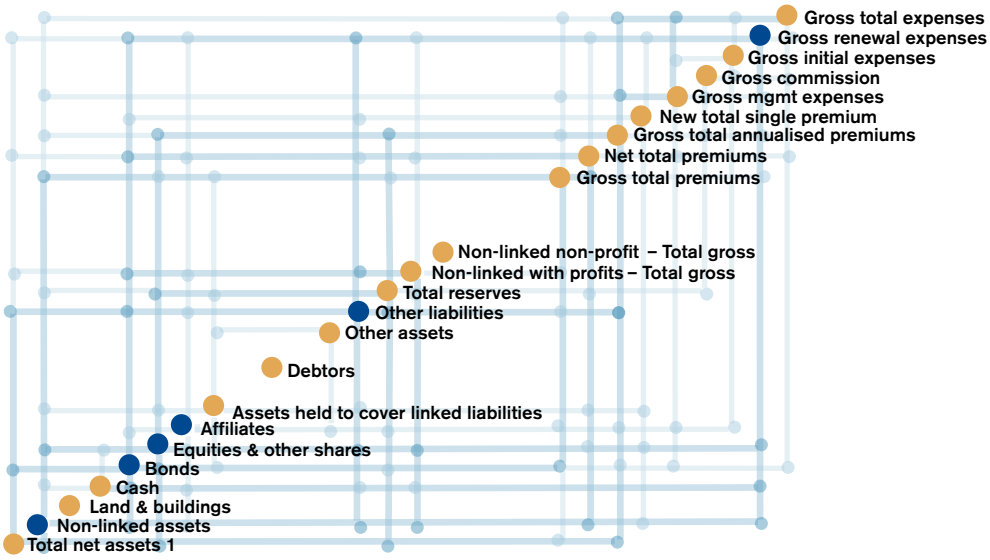

Source: Milliman analysis on FSA returns data (SynThesys) using DACORD (DRTS)

At this stage some of the asset variables, other liabilities, and renewal expenses show up as dominating the picture as we calculate that their influence on the other variables is most significant. As we move ahead to 2001, in Figure 5, we can see that the impact on premiums is beginning to bite as their role in the structure of the sector's performance becomes more dominant, and reserves similarly become more important.

FIGURE 5: INFORMATION SHARING BETWEEN KEY INDICATORS IN 2001

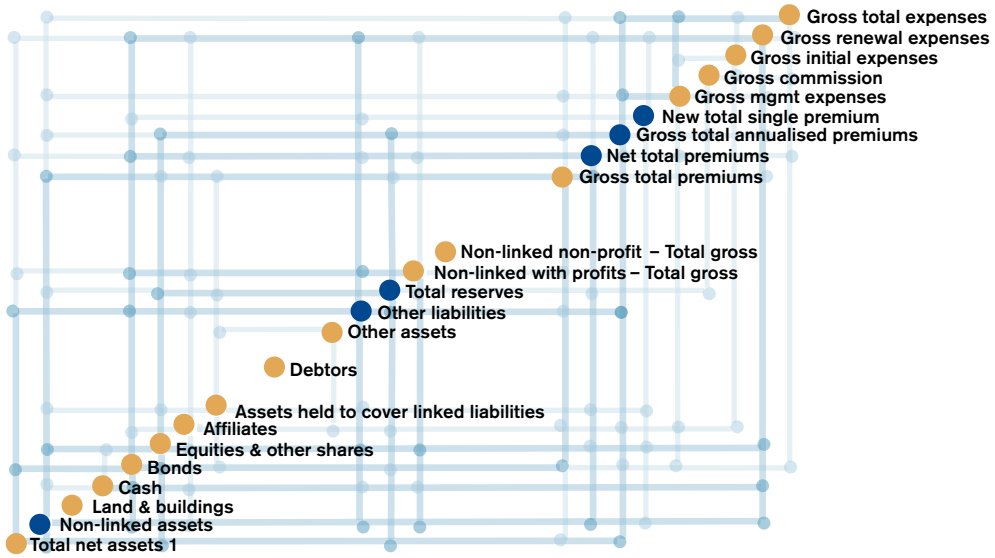

Source: Milliman analysis on FSA returns data (SynThesys) using DACORD (DRTS)

If we cycle ahead to 2004, after the recovery has been underway for a year, the industry reaches a peak in terms of the complexity of how information is shared between indicators. It is pertinent that at this stage a strong connection appears between debtors and reserves and slightly weaker connections to a number of indicators. Prior to 2004, debtors played a limited role in the landscape, but thereafter is a persistent feature. While it is not a dominant indicator it is an interesting feature of the post-2003 growth period, particularly given that commentators on the current crisis cite leverage as one of the key precipitators.

Following 2004, we see a steep decline in the complexity of how the indicators are connected. This may be in part due to a focus on consolidation, rebuilding, and so on. For 2007, we see that the most important indicator is reserves. Interestingly, since 2002 total reserves have shown a strong connection to land and buildings assets, indicative of the importance of property in the UK investment landscape.

Since 2002 total reserves have shown a strong connection to land and buildings assets, indicative of the importance of property in the UK investment landscape.

DOES SIZE MATTER?

In the previous section we looked at the behaviours of the life assurance sector as a whole, but this section looks at whether the size of an organisation matters in looking at its vulnerability to performance issues.

The chart in Figure 6 shows that there are some differences between the trends in the uncertainty of the different size groupings. The industry as a whole has experienced slightly increased uncertainty from 2000 to 2006 with a drop in 2007 as a result of the recent collapse in markets. Figure 6, however, shows that not all size groups have followed this trend.

FIGURE 6: RELATIVE UNCERTAINTY BY COMPANY SIZE

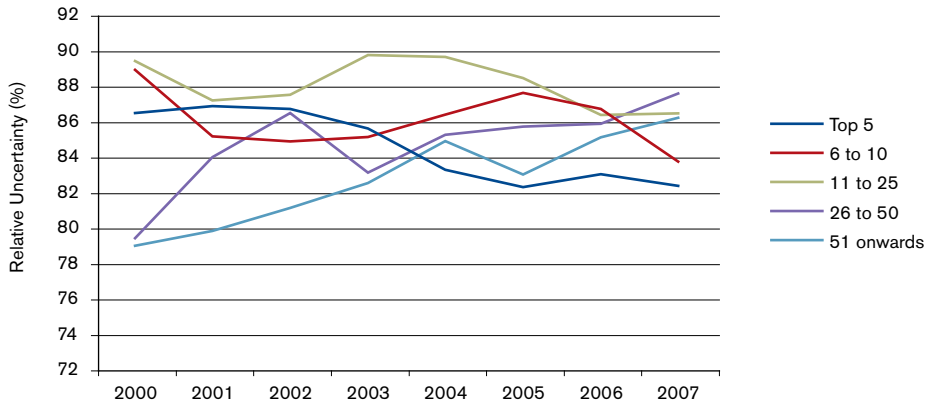

Source: Milliman analysis on FSA returns data (SynThesys) using DACORD (DRTS)

The top five companies in terms of size have generally decreased their uncertainty. It is actually a relatively large drop and may indicate the difficulties these companies have in generating high performance. The smallest companies have significantly increased their uncertainty. Although not operating quite at dangerous levels yet, the trajectory suggests that performance is becoming more fragile. The mid-tier (11–25) companies initially grew more uncertain and reached a dangerous level around 2003 after which their uncertainty has steadily drained back to more sustainable levels. This tends to suggest that this group suffered more than the others during the 2000–2003 crash.

The higher level of uncertainty for smaller companies may reflect the fact that their strategies are narrower and more optimised, but the fact that the level of uncertainty is consistently growing is not a good story.

The higher level of uncertainty for smaller companies may reflect the fact that their strategies are narrower and more optimised, but the fact that the level of uncertainty is consistently growing is not a good story. So, based on the evidence of Figure 6, it seems that there may be trouble ahead for smaller companies, notwithstanding the current economic situation, if this build-up in uncertainty gets beyond a level that they can handle.

If we look at the distribution of uncertainty by size at a single point in time, we find that the distribution is broadly uniform, with a slight tendency towards higher uncertainty among the smaller companies. Although the distribution of uncertainty is quite uniform, we do observe a tendency for smaller companies to exhibit a wider variance in uncertainty. This is clearly visible in the chart in Figure 7.

FIGURE 7: VARIANCE IN RELATIVE UNCERTAINTY BY COMPANY SIZE

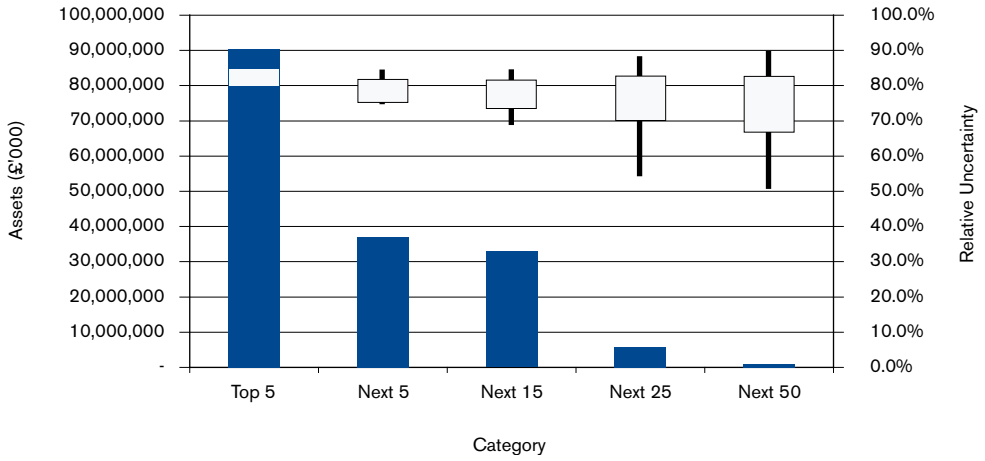

Source: Milliman analysis on FSA returns data 2007 (SynThesys) using DACORD (DRTS)

SHAREHOLDERS VERSUS MEMBERS

The next question we addressed in this study was whether there were differences in the performance of proprietary companies and those owned by members (mutuals, friendly societies, cooperatives, etc.). For the purpose of this analysis we have assumed that companies were always of the type that they are in 2007, i.e., demutualised companies are classified as proprietary throughout. This enables us to track the current cohort of mutuals rather than just the sector.

Overall the uncertainty of proprietary and mutual companies as groups seems to be out of phase.

Overview

Overall the uncertainty of proprietary and mutual companies as groups seems to be out of phase. In broad terms, it seems that the mutuals suffered a significant drop in uncertainty from 2000 to 2001 whereas proprietary companies increased. Both groups seemed to pause as the 2001–2003 turmoil played out, and then the uncertainty level of mutuals steadily climbed as proprietary companies fell. This can be seen in Figure 8.

FIGURE 8: RELATIVE UNCERTAINTY FOR PROPRIETARY AND MUTUAL OWNERSHIP

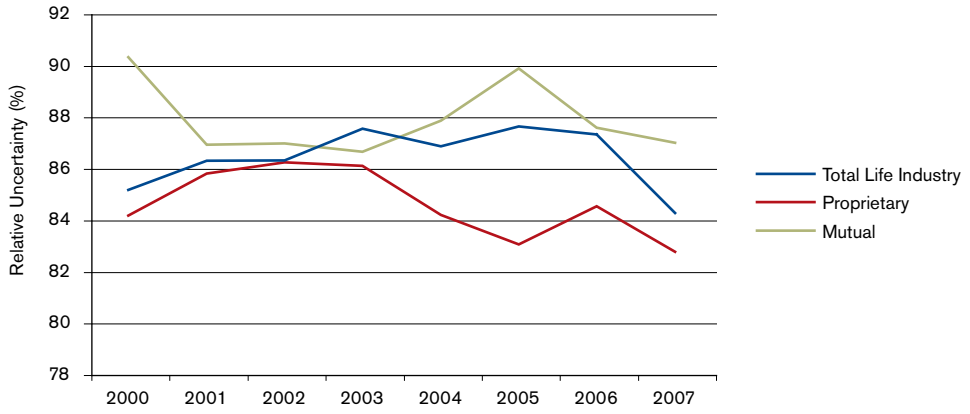

Source: Milliman analysis on FSA returns data (SynThesys) using DACORD (DRTS)

When we look at the complexity of the two ownership types, allowing for the structure of the performance variables in addition to uncertainty, we see that this also diverges, with mutuals generally decreasing and proprietary companies generally increasing, as shown in Figure 9. The fact that the uncertainty measure for proprietary companies decreases as overall complexity increases suggests that the increase in complexity is coming from the structure of the performance, through an increase in rigidity/interconnectivity.

FIGURE 9: COMPLEXITY FOR PROPRIETARY AND MUTUAL OWNERSHIP

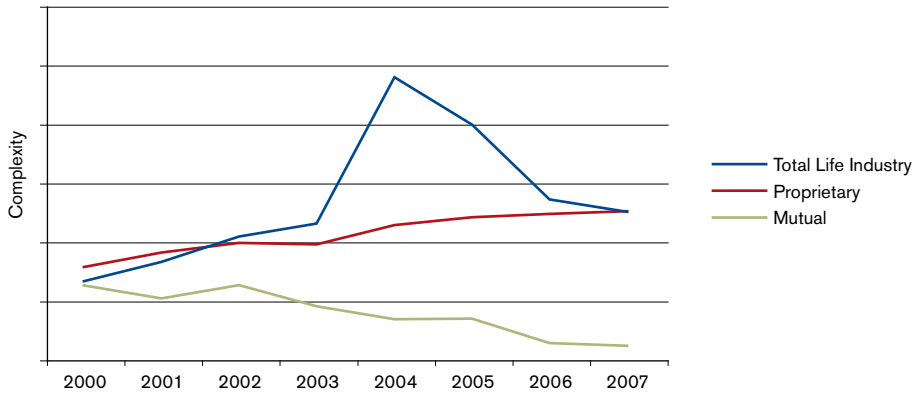

Source: Milliman analysis on FSA returns data (SynThesys) using DACORD (DRTS)

Overall we can also see in Figure 8 that the level of uncertainty for mutuals seems to run a bit higher than for proprietary companies. This may reflect their more focused strategies and hence a degree of fragility arising from specialisation. Their lower complexity therefore arises from a lower rigidity compared to proprietary companies.

The level of uncertainty for mutuals seems to run a bit higher than for proprietary companies. This may reflect their more focused strategies and hence a degree of fragility arising from specialisation.

FIGURE 10: COMPARING DISTRIBUTION OF RELATIVE UNCERTAINTY BY SIZE FOR MUTUAL AND PROPRIETARY COMPANIES

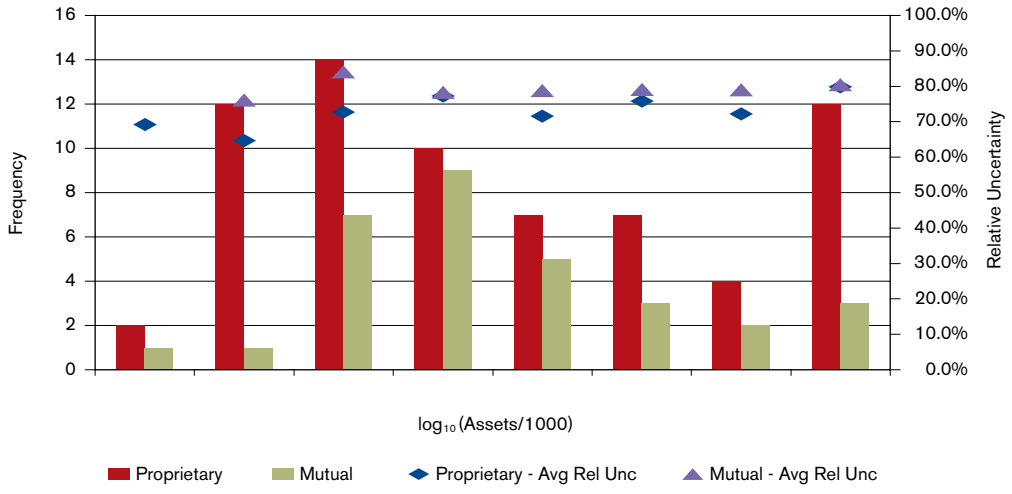

Source: Milliman analysis on FSA returns data 2007 (SynThesys) using DACORD (DRTS)

At each size category we also see that mutuals tend to run at a slightly higher uncertainty than the equivalent proprietary companies on average (see Figure 10). However, the smaller proprietary companies have a higher dispersion around the average.

Proprietary Companies

The proprietary companies are the majority of companies in the sector and account for much of its assets. So a hypothesis is that they tend to follow a similar trend to the overall sector. This is broadly true for larger companies that currently tend to be proprietary, while the behaviour of the smaller companies is more influenced by mutuals.

So the larger firms show moderate declines in uncertainty and modest fluctuations in complexity. The mid-tier firms show high but slowly falling uncertainty, and relatively modest complexity. And the smaller firms show increasing uncertainty and larger, more volatile, complexity.

So the larger firms show moderate declines in uncertainty and modest fluctuations in complexity. The mid-tier firms show high but slowly falling uncertainty, and relatively modest complexity. And the smaller firms show increasing uncertainty and larger, more volatile, complexity. These trends in uncertainty and complexity can be seen in Figures 11 and 12 respectively.

FIGURE 11: RELATIVE UNCERTAINTY FOR PROPRIETARY COMPANIES

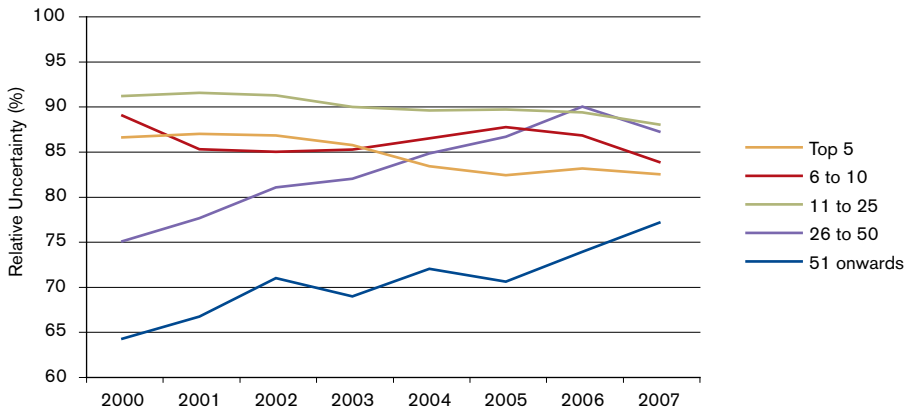

Source: Milliman analysis on FSA returns data (SynThesys) using DACORD (DRTS)

FIGURE 12: COMPLEXITY FOR PROPRIETARY COMPANIES

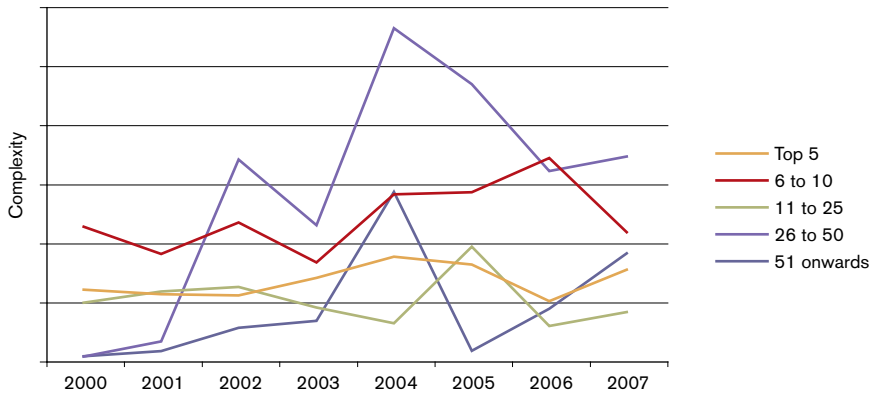

Source: Milliman analysis on FSA returns data (SynThesys) using DACORD (DRTS)

Proprietary companies show that uncertainty generally increases by size, but with a tendency for smaller companies to have more variation around the average, as shown in Figures 13 and 14.

FIGURE 13: RELATIVE UNCERTAINTY BY COMPANY SIZE FOR PROPRIETARY COMPANIES

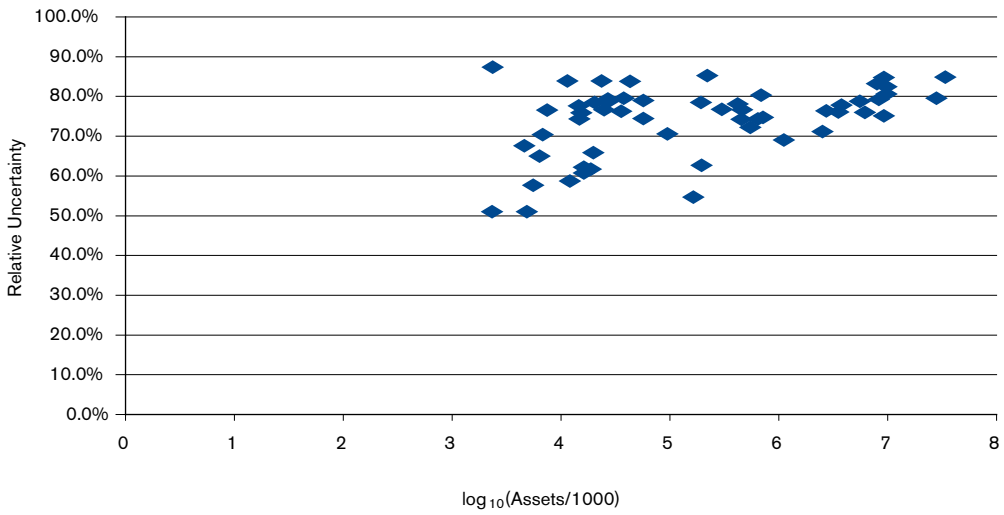

Source: Milliman analysis on FSA returns data 2007 (SynThesys) using DACORD (DRTS)

FIGURE 14: VARIANCE IN RELATIVE UNCERTAINTY FOR PROPRIETARY COMPANIES

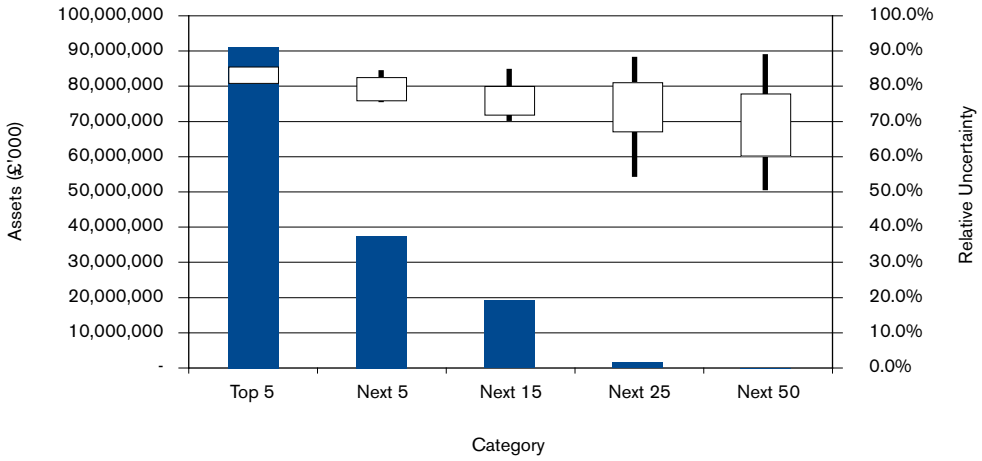

Source: Milliman analysis on FSA returns data 2007 (SynThesys) using DACORD (DRTS)

Mutual Companies

The mutual companies tend to be smaller than proprietary companies, although the largest are in the overall top 20. So the behaviour of the industry's medium to smaller-sized companies is more influenced by mutuals than the top 10.

The uncertainty levels of the larger mutuals have been falling before stabilising around the end of the 2000–2003 crash.

The uncertainty levels of the larger mutuals have been falling before stabilising around the end of the 2000–2003 crash. Levels rose again until about 2005 before the largest mutuals remained more constant, while the second group fell rapidly. The smallest mutuals have been gradually increasing in uncertainty. Figure 15 shows the trends in uncertainty.

FIGURE 15: RELATIVE UNCERTAINTY FOR MUTUAL COMPANIES

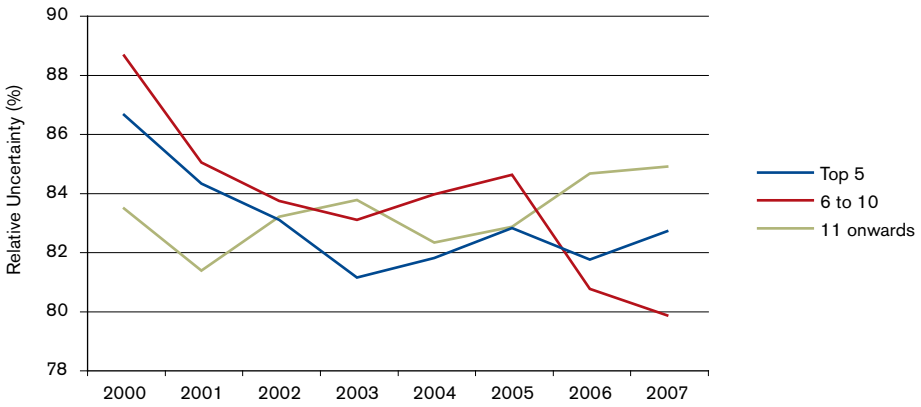

Source: Milliman analysis on FSA returns data (SynThesys) using DACORD (DRTS)

In terms of complexity, the mid-tier mutuals seem to have consistently higher complexity, but all mutuals have generally been reducing complexity. This is clearly in contrast to the proprietary companies. We observe that the fall in complexity for the large mutuals appears to be driven by a reduction in the rigidity of the organisations. This is in contrast to the largest proprietary companies whose rigidity has somewhat increased.

Mutuals show a slightly more uniform distribution of uncertainty by size than proprietary companies. This is shown in Figure 16.

FIGURE 16: RELATIVE UNCERTAINTY BY COMPANY SIZE FOR MUTUAL COMPANIES

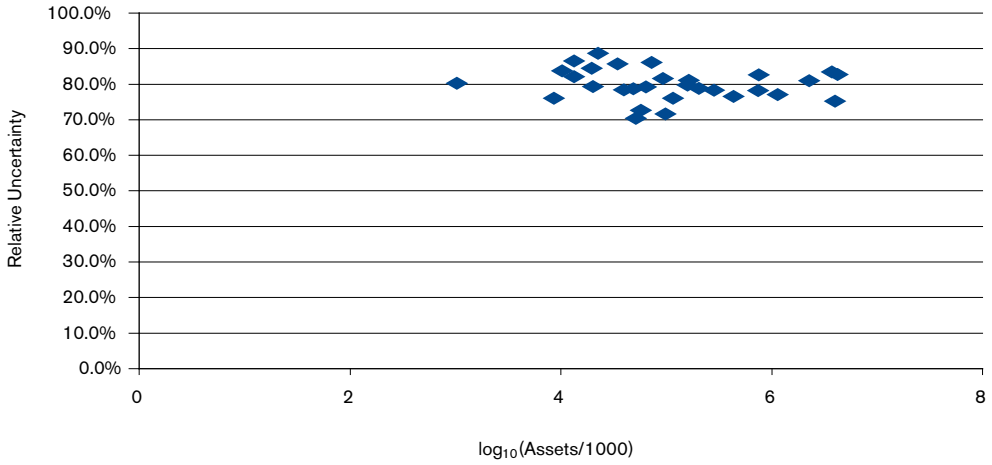

Source: Milliman analysis on FSA returns data 2007 (SynThesys) using DACORD (DRTS)

The variance in uncertainty at smaller sizes is slightly higher than for larger mutuals but the average varies less by size than for proprietary companies, as shown in Figure 17.

FIGURE 17: VARIANCE IN RELATIVE UNCERTAINTY FOR MUTUAL COMPANIES

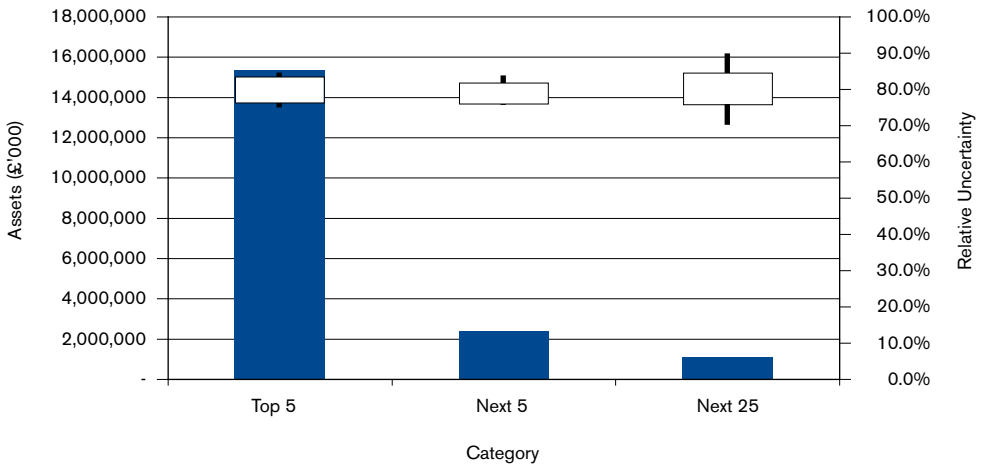

Source: Milliman analysis on FSA returns data 2007 (SynThesys) using DACORD (DRTS)

CONCLUSIONS

Through our analysis of the UK life insurance sector, we have shown that headline performance indicators do not always tell the whole story. Ultimately, these indicators will reflect poor results, but they require time to integrate bad news. By using a combination of complexity-based tools it is possible to look deeper into the information about the sector's performance and seek out the patterns that truly signify its current state.

This report shows that the signs of recent crises could have been detected even from relatively high-level information. It also shows that not all companies are homogeneously impacted by chains of events.

This report shows that the signs of recent crises could have been detected even from relatively high-level information. It also shows that not all companies are homogeneously impacted by chains of events. Even in the face of extreme systemic risk some companies are clearly more affected than others.

Our use of complexity-based techniques uncovers the true underlying features of what is going on and enables us to better identify the onset of trouble. We have identified and analysed interesting features of performance for various company types. Indeed, these methods are applicable to any industry. It is clear that not all performance issues that will arise come from the present economic situation—their genesis lies elsewhere and is more fundamental.

Milliman, whose corporate offices are in Seattle, serves the full spectrum of business, financial, government, and union organizations. Founded in 1947 as Milliman & Robertson, the company has 49 offices in principal cities in the United States and worldwide. Milliman employs more than 2,100 people, including a professional staff of more than 1,100 qualified consultants and actuaries. The firm has consulting practices in employee benefits, healthcare, life insurance/financial services, and property and casualty insurance. Milliman's employee benefits practice is a member of Abelica Global, an international organization of independent consulting firms serving clients around the globe. For further information visit www.milliman.com.

Neil Allan
University of Bristol
N.D.Allan@bath.ac.uk

Neil Cantle
neil.cantle@milliman.com

Finsbury Tower,
103-105 Bunhill Row
London, EC1Y 8LZ
+44 20 7847 1500

milliman.co.uk