MILLIMAN RESEARCH REPORT

Lapses in concentration

June 2017

Jennifer Smith, FIA
Neil Cantle, MA, ASA, FIA

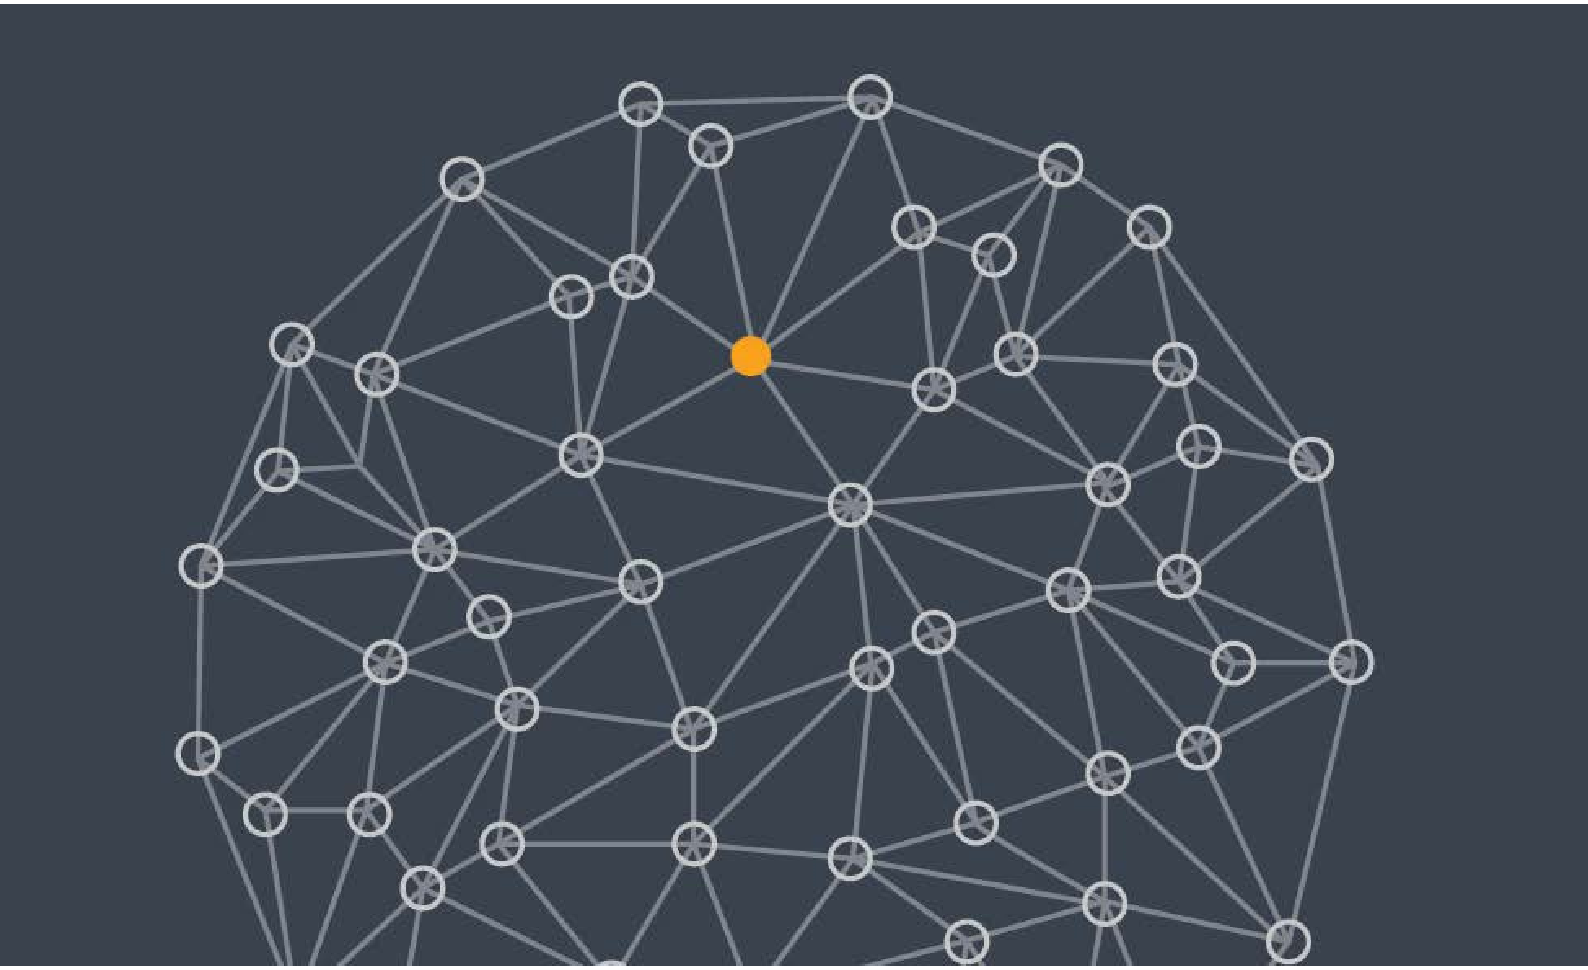

Table of Contents

1. EXECUTIVE SUMMARY

1.1. The purpose of this report is to set out the methodology and results of Milliman's research investigation into the use of advanced systems mining techniques to determine how lapse experience for long-term insurance business might change according to the prevailing dynamics within the business and due to uncontrollable external factors.

1.2. The research was carried out in conjunction with a client of the London office. The client provided the relevant data and helped to explain potential underlying meaning for the results obtained.

BACKGROUND

1.3. Milliman has previously worked on client projects where the aim was to model base lapse and mass lapse rates and to investigate to see if the drivers varied between the two scenarios. Model structures were developed using a cognitive analysis of expert opinions about the dynamics of lapses. The work reveals that a range of factors interact to influence lapse behaviour. And yet this is typically not taken into account directly when setting assumptions.

1.4. Actuaries are often called upon to make long-term assumptions for lapses, and these are often set with reference to historical data about similar products and sales channels.

1.5. It is typically assumed that policyholder behaviour is influenced strongly by duration in force—the argument usually being that this relates to upfront penalty periods—or some other simple factor which the policyholder rationally takes into account when deciding to lapse. But is this actually true?

1.6. To investigate this, we decided to analyse some time series data of lapse rates provided by a client, alongside some internal performance metrics and external market data, to see if there were patterns which could help to identify trends in lapse rates.

1.7. The client kindly offered to assist us with our investigation and provide data and supporting contextual commentary.

1.8. The data was analysed using DACORD[1], a tool developed by DRTS which enables users to analyse multiple time series of data and uses information metrics (Shannon entropy and graphical complexity) to look for non-linear relationships between the variables.

KEY RESEARCH OBJECTIVES

1.9. We set out to investigate two key areas, based upon the challenges commonly faced by our clients:

- To assess whether drivers of lapse rates vary over time, and to what extent changes in the underlying drivers can be identified or predicted.
- To make use of a variety of indicators and data to reflect that individuals sometimes behave irrationally or are motivated by seemingly unrelated factors. The aim is to identify drivers of lapses beyond the typical ones (e.g. duration in force).

SCOPE AND STRUCTURE OF THE REPORT

1.10. In Section 2, we introduce the data that we collected to perform the research, explain where we sourced it from and the practical problems that we encountered with this process.

1.11. Then, in Section 3, we explain the process used to perform the research and describe the techniques used to analyse the data to look for drivers of lapse and the trends in these over time.

1.12. In Section 4, we explore the results provided by the analysis for the three different product types, before exploring, in Section 5, how these could be used or further developed in the future.

[1] https://www.dacord.co.uk/

KEY FINDINGS

1.13. In Section 6, we conclude with two key findings that form the core results of the research.

1.14. First, the drivers of lapse rates vary by line of business and depend on how the policyholder uses the policy and whether the policyholder is an individual or corporate client.

1.15. Second, the drivers of lapse rates do change over time and can build to a tipping point where circumstances can change very suddenly. The rate of lapse for a particular product may, therefore, appear to remain constant over a lengthy period of time when the underlying drivers are evolving. At some point, the new drivers may lead to a tipping point or other unexpected outcome.

1.16. In setting long-term assumptions for lapse, the research suggests that there is merit in thinking more deeply about whether past drivers can reasonably be expected to persist. This is particularly helpful when thinking about drivers for new products, where there is no past data that you can use to support an assumption and where drivers may differ from existing products. It is also highly relevant for mass lapses, where the key drivers tend to be something new which is not present in past data or where a driver’s influence has significantly changed. Collecting information about the drivers of lapse behaviour also offers an opportunity to identify emerging trends which might be useful inputs to servicing and retention strategies.

2. DATA COLLECTION

2.1. The first stage of the project was to collect the data to be analysed.

2.2. There were two primary forms of data that we were interested in collecting: internal data provided by the client detailing the lapse rates and other corresponding internal performance indicators, and external data we could research ourselves from publically available data sources and use to add context to the internal data.

INTERNAL DATA

2.3. The client is a wealth manager and so operates in a slightly different way than most traditional life insurance companies. It offers its products through financial advisers who advise their clients on a wide range of financial solutions from the client, although they can advise on products from other providers if required. Such products include bonds, pensions, unit trusts and ISAs.[2]

2.4. The client's target customer is someone of 'mass affluent' wealth, who would typically have £50,000 or more of assets available to invest. Retention rates are generally high, which reflects the strong bonds that clients have with their adviser.

2.5. We asked the client to provide us with as much time series data as possible concerning its internal metrics that may be linked to lapse rates. Even if a metric did not seem to obviously link to lapse rates—for example, policy age at entry—we wanted to analyse the data anyway to see if the data field might be a proxy for some underlying behaviours which might drive the lapse rates.

2.6. The initial set of data that we received was a set of 'movement' files citing individual policy movements, including deaths, lapses and partial lapses. To convert the data into the form we could use for our analysis, namely monthly lapse rates, proved challenging, but this was soon resolved by clarifying exactly which fields referred to the data we needed.

2.7. Following on from the lapse data, the client started to ask various internal departments about other available data sources that we could use. To minimise the impact on the client, we were interested particularly in data that the company already collected via its computer systems or reporting that we could use for an alternative purpose.

2.8. One area in particular where we obtained interesting and informative data was from the services team. It provided time series data covering:

- Complaints
- Call data such as waiting and hold times
- Written enquiries
- Compensation paid

2.9. We were also able to obtain some more detailed information about the policyholders whose lapses we were analysing, such as age at entry and duration in force, to add to the analysis.

2.10. The data provided by the client covered three main product types: ISAs, personal unit trusts and corporate unit trusts. To try and give the most granular analysis possible, we decided to analyse the three product lines separately, in anticipation that different product lines might exhibit different behaviours and drivers for lapsing.

[2] Individual Savings Accounts (ISAs) are a type of UK investment vehicle which UK residents can invest in without incurring taxation on investment income or capital gains. There is an annual limit (currently £20,000) on the amount that can be invested and a restriction on the types of asset that can be invest in.

EXTERNAL DATA

2.11. External time series data for various indicators was collected from publically available data sources to explore how lapse rates behaved under different financial and industrial conditions.

2.12. The anticipation was that indicators should not be limited to those with a direct, or widely accepted, link to lapses. Instead, many types of data were collected with the aim of identifying underlying factors driving lapses. For example, data on levels of international travel would not generally be considered as being related to lapse rates, but the underlying driver of spending capacity and a person's desire to spend on non-essential items may suggest a link between the two variables.

2.13. Various market and demographic data metrics were obtained from the Office for National Statistics (ONS) website,[3] which provides a comprehensive range of freely available time series data for economic, industry, employment and population metrics. The choice of this data source reflects the quality of the data and the availability of frequent (i.e., monthly) statistics covering a long time period in the majority of cases. Statistics sourced from the ONS include:

- Unemployment and redundancy rates
- Social indices including spending on travel, leisure and entertaining
- Inflation
- House price indices

2.14. Other sources of external data used included:

- FTSE 100 index values
- IMF UK market indicators, including metrics such as gross domestic product (GDP), government revenue and expenditure and investment as a percentage of GDP
- World Bank UK market indicators

PROCESSING THE DATA

2.15. A key decision during the set-up phase was the choice of data frequency. Technically, an annual system is different to a monthly or quarterly one, and the interpretation of results needs to be done in that context. Due to the relatively few years of data for some variables, we chose to look at the 'monthly' system and therefore only used variables which had observations at this level of granularity.

2.16. The analysis did work well with the monthly data, as this allowed us to spot the trends on a much more granular level and gave us a better understanding of the changes between the data points. In particular, it enabled us to explore trends which could be potential precursors to lapse behaviour which may have been harder to spot when aggregated to an annual observation.

LAGGING THE DATA

2.17. The final data process that we carried out was to try and time-lag the data. Once we had established the data items that we were to analyse, we started to consider how these data items might influence policyholders and hence make them lapse. For certain data items, it soon became clear that such items might have a delayed effect on a person's decision to lapse.

2.18. For example, if a policyholder received poor customer service on a call to ask a query about their policy, it is likely that they would take a small amount of time to consider their options before lapsing. This might mean that some months may pass before the policyholder lapses a policy, despite the driver of that lapse occurring months before.

2.19. Because of this thinking, we decided to carry out additional investigations with the research to see what happened when we lagged the data. We first lagged some variables by one month, to capture those drivers where decisions are likely to be delayed but relatively immediate, and second by three months, to capture the situations where the policyholder would need to consider the decision more carefully and where it may take time to put alternative arrangements in place.

[3] Office for National Statistics, https://www.ons.gov.uk/

3. PROCESS

IMPORTING THE DATA

3.1. Once the data had been collected and processed into an appropriate format, the next step was to import the data into DACORD and start to form the analysis.

3.2. At this stage, we had nine sets of data covering three different products and including no lag, a lag of one month and a lag of three months, and intended to carry out separate analysis on each set of data. As a future enhancement, we could consider lagging variables differently to reflect the time each one takes to influence the outcome.

3.3. The process of uploading data to DACORD is very simple and involves uploading a CSV file. DACORD will verify to the user if the file is suitable for use, and once this process is complete, the analysis can begin.

ANALYSIS

3.4. DACORD provides a range of different views on the imported data, which allows the user to look in different ways for patterns or trends which may be present.

3.5. The process uses a user-specified number of time steps as a learning period to non-parametrically estimate the probabilities required for the information, mutual information and graphical complexity calculations. These are used to reveal non-linear interactions between the data and provide metrics relating to the underlying system structure as it evolves over time. For this process, a larger window is better, as it enables the algorithms to be more certain about the parameter estimates and reduces the volatility of the results. The window for our analysis was set at 12 data points (i.e., one year), and the end of the learning period can be seen on some of the graphs below as a red vertical line.

ORIGINAL DATA

3.6. The first piece of analysis we carried out was to simply observe the original data in a time series format and look for any obvious visual outliers. We normalised the data set to make it easier to visually compare all the data fields on one graph.

3.7. The graph below shows the set of normalised data points for the personal unit trust business, without any time-lags in place.

FIGURE 3.1: NORMALISED DATA

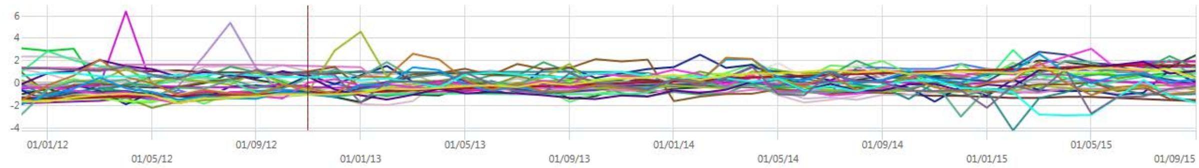

3.8. This view helps to identify anomalies in the data, which in turn provides follow-up questions to the client as to why these anomalies occurred and whether there may be errors in the data. The graph below in red circles highlights potential areas of further investigation.

FIGURE 3.2: NORMALISED DATA WITH AREAS OF FURTHER INVESTIGATION

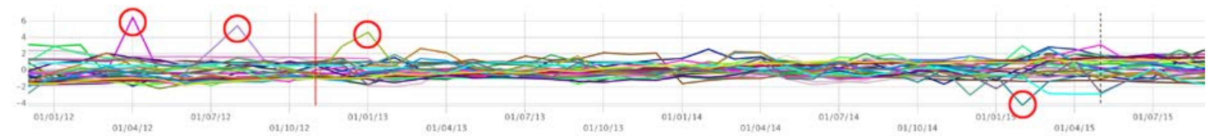

GRAPHICAL COMPLEXITY

3.9. Another method of analysing the data is to look at the graphical complexity of the system, which can be seen in the graph below.

FIGURE 3.3: GRAPHICAL COMPLEXITY

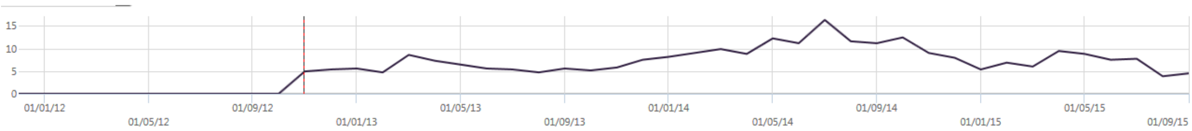

3.10. Graphical complexity is essentially a measure of how difficult it is to construct a particular system. The more elements there are and the more complicated their relationships, the harder it would be to create and hence can be described as being more 'complex.' A more complex network can be prone to 'tipping points' and other non-linear behaviours; these are interesting to explore, as they can precede something unexpected, like a mass lapse event, for example.

TOTAL UNCERTAINTY

3.11. Another view of the data is through the lens of uncertainty. This measure is scaled as a percentage of the maximum possible uncertainty (i.e., where all outcomes in the window are equally likely to occur). When a system's behaviour is entirely predictable, the uncertainty is 0%. When the system is as unpredictable as possible, the uncertainty is 100%. As can be seen in Figure 3.4, the lapse 'system' tends to exhibit moderately high levels of uncertainty a lot of the time. This is not unexpected for a financial system, so we will be more interested in variations and transitions.

FIGURE 3.4: TOTAL UNCERTAINTY

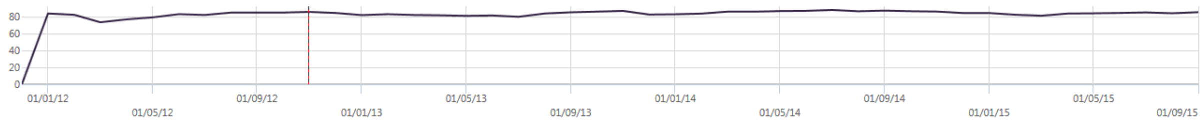

3.12. When a system exhibits sudden changes in uncertainty, this can be indicative of the system being under stress and about to exhibit some behaviour which 'releases' that and returns the system to a more predictable state. We are therefore interested to know if such signals occur prior to a lapse event.

COMPARE SYSTEMS

3.13. It is also useful to combine the above measures by studying a plot of total uncertainty against graphical complexity. The path plotted on the graph over time is a way of uncovering the way in which the system is behaving, particularly in relation to periods of stress. A path which takes the system towards the top right indicates that the system becomes more fragile, by virtue of being more complex, and more stressed by virtue of showing greater uncertainty. When the system suddenly moves down and left away from that region, we might infer that some kind of release has taken place: A tipping point has been reached.

3.14. In the following example graph, the potential tipping point can be identified in red. Each point on the graph represents a point in time, and it is possible to watch how the complexity and uncertainty develops on the graph over time. The point highlighted in red represents August 2014, and so this provided a point of further investigation in the analysis allowing us to see if the system substantially changed after this point.

FIGURE 3.5: IDENTIFYING TIPPING POINTS

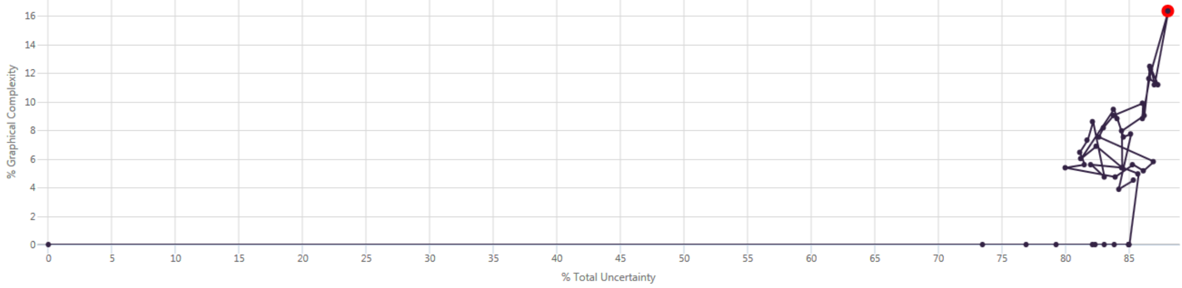

3.15. In practical terms, a system changing after a tipping point could mean that the drivers of lapses changed, which would be very helpful information in the investigation of drivers of lapse rates.

EXPLORE

3.16. Another interesting view of that data is to lay out the variables as 'parallel co-ordinates' and join the values they each take at a point in time with a line. This provides a simple way to visualise multi-variate data—axes can be moved around next to other variables to more easily see they joint behaviours, and some variables can be constrained to see which values of other variables are consistent with this.

3.17. Another benefit of this view is that it is easy to spot outliers, as they are much more visually obvious. The sixth axis on the example below (UT Non-Corp – Number) seems to have one value which is very different to all of its other values, for example.

3.18. In the graph below, the red line shows all the data points that were observed in November 2011. It is then possible to consider the movement of this to see how a particular data field changes over time.

FIGURE 3.6: DACORD'S EXPLORE VIEW

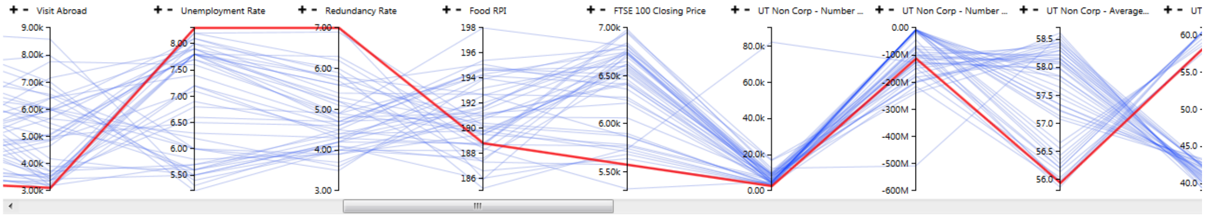

Date	Jewellery	Perfume and···	Restaurants ···	Creative Arts	Accomodati···	Travel Agency	Sports Activities	House Prices	Visit Abroad	Unemploym···	Redundancy···	Food RPI	FTSE 100 Clo···	UT Non Cor···	UT Non Cor···
01/12/11	142.5	100.6	91.7	8.1	98	3	95	173	3113	8.3	7	188.8	5572.299805	2646	-112995152.1
01/01/12	141.4	100.3	91.9	5.6	99.4	1.7	95.7	175.5	3260	8.3	7	186	5681.600098	3227	-215385633.4
01/02/12	142.3	100.5	92.1	1.9	97.6	5.1	94.7	172.8	3364	8.2	6.9	190.2	5871.5	3848	-177082130.7
01/03/12	124.7	100.4	92.3	2.8	100.4	2.9	94.4	172.6	3635	8.2	6.2	188.7	5768.5	5907	-147923610.9
01/04/12	123.6	100.2	93	2.2	100	-4.8	95.2	175.1	4845	8.1	5.9	185.2	5737.799805	82172	-128032703.6
01/05/12	122.7	99.8	93.2	4.8	99.7	-4.4	94.5	175.1	4658	8	6	188.4	5320.899902	5027	-43867924.46
01/06/12	122.8	100.6	93.5	5.3	99.6	-1.5	94.4	177.1	6018	8.1	5.7	187.3	5571.200195	3742	-71548383.69
01/07/12	122.6	100	93.7	9.4	100.8	-5	116.7	179.1	5662	7.9	5.2	186.6	5635.299805	3471	-89229456.37
01/08/12	122.7	100.2	93.8	9.7	103.1	-0.6	140.8	179.3	7158	7.9	5.1	189.8	5711.5	3672	-67835026.97
01/09/12	123.7	100.5	93.9	8.7	101.6	1.2	107.4	178.2	6399	7.9	5.8	190.1	5742.100098	3199	-11455101.73
01/10/12	124.1	100.6	93.9	2.5	100.6	4.5	89.3	176.7	5095	7.8	6.3	189.5	5782.700195	6822	-13586494.26
01/11/12	123.6	100.5	94.3	-2.4	100	1.3	89.7	177.6	3368	7.8	5.8	192.7	5866.799805	3865	-6448570.285

3.19. This analysis can be played across a period, where each time period (in our case, a month) is represented by a red line on the graph. In playing the analysis through, it is easier to spot sudden jumps in the data points, which might highlight an area for further investigation.

3.20. This feature also allows relationships between different variables to be identified through patterns in the analysis. For example, it might be possible to conclude that as one variable increases, the other decreases—shown by parallel lines going from the high of one graph to the low of the other.

3.21. It is possible to filter this analysis by the range of data points. Suppose, for example, that we were only interested in analysing the system when the FTSE 100 was above 6500. We can then crop the analysis such that it only looks at time periods where the FTSE 100 was greater than 6500, as shown in the graph below.

FIGURE 3.7: NARROWING THE RANGE IN EXPLORE

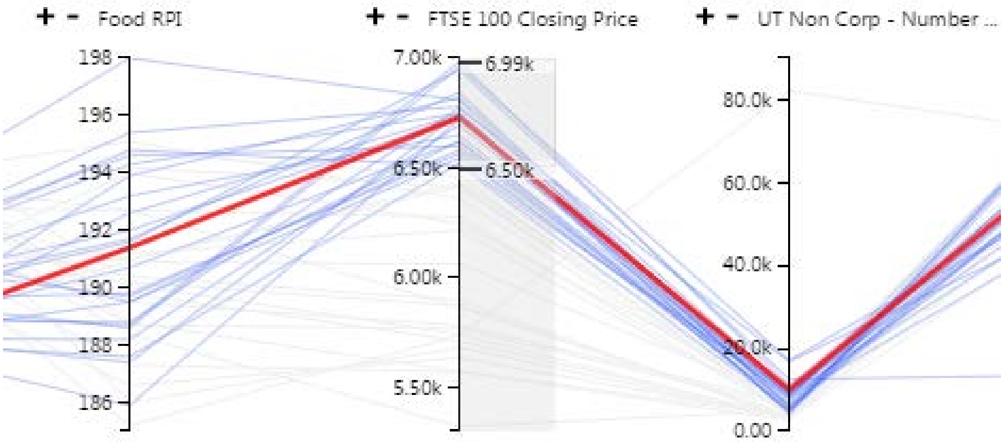

3.22. This feature also allows the user to drag and drop the data fields around the analysis for ease of comparison between data fields.

4. RESULTS

4.1. The following section outlines the results that were drawn from the lapse rate research.

BETWEEN LINES OF BUSINESS

4.2. An interesting finding of the research was that the drivers of lapse varied by line of business. The sections below explain the drivers per line of business in more detail.

ISA

4.3. The ISA business had very few links to any of the drivers in our analysis, and the links were not consistent over time. This is shown in the following graph, which plots the link strength for the number of lapses on ISA products for the period studied.

FIGURE 4.1: LINK STRENGTH OF THE ISA LAPSE RATE

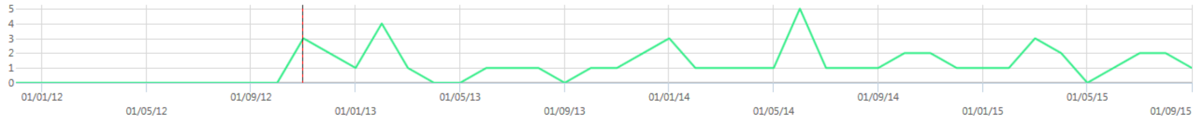

4.4. The y-axis shows the number of links that the lapse rates have to the system. As you can see from the graph, the number of links to the system never exceeds 5 and most of the time is in the 1-2 range. As we analysed the Proportion of Units Lapsed which was equal to the Number of Lapses divided by the Total Number of Units, we would typically expect to see links to those variables anyway, which makes the lapses observed in this analysis even less significant.

4.5. This suggests that people withdraw money from their ISAs for a wide range of reasons and that periodically some of those included the drivers in our analysis. This makes logical sense, as people often withdraw money from their ISAs for very personal reasons—for example, to purchase a car, to start up a business or to pay for house renovations. We might therefore suppose that a wider set of external data relating to the customer's world would be useful for predicting lapses for this type of product.

UNIT TRUST (PERSONAL)

4.6. The primary drivers of the unit trust business for personal sales were drivers to do with the service levels of the company, such as complaints and service times. This is reasonable given that the client's customers are generally very loyal to their adviser. Generally, it would only be something related to service that would cause a customer to lapse a policy, with some sort of breakdown in communication being at the core of the problem.

4.7. The behaviour of personal clients appears slightly more emotionally driven than that of corporate clients. Individuals want to be treated with respect and understanding, and an adverse experience can significantly influence their perception of the suitability of the product and the company for them.

4.8. The only exception to the service factors that were observed to be drivers of the business was unemployment. However, this does fit with the concept that the personal customers are affected by what directly affects them, as opposed to the drivers that affect the market or the product itself.

4.9. The image below shows the strength of the links for the lapse system of the personal unit trust business as at June 2014. As described, the key drivers are service levels—for example, breaches received and number of telephone calls—as well as personal drivers such as unemployment.

4.10. Links are shown as lines between concepts, where strong links are red, medium links are orange and weaker links are blue. Each variable is represented by a dot in the graph, and a graph such as the one below is produced for each of the time periods that are analysed. These can be played sequentially, a bit like watching a movie, so it is possible to spot trends emerging or unusual behaviour in a particular time period.

FIGURE 4.2: LINKS TO LAPSE RATES IN THE PERSONAL UNIT TRUST BUSINESS

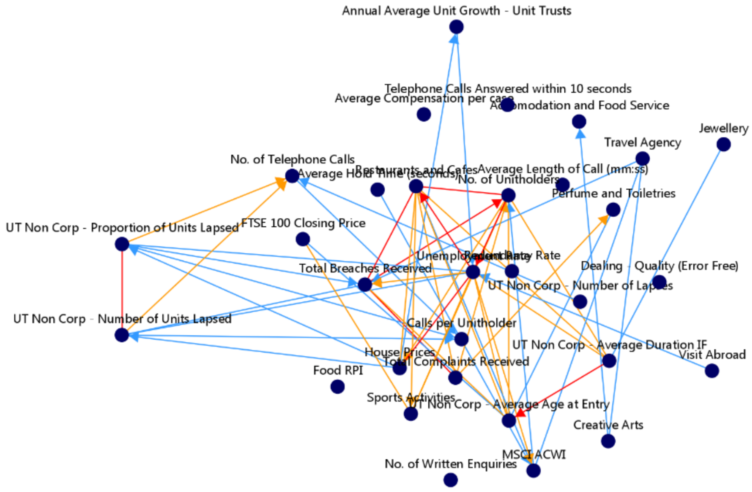

UNIT TRUST (CORPORATE)

4.11. The primary drivers of the corporate business were more financially driven than the personal business, and included drivers such as the FTSE100 index, redundancy rates and house price indices. This suggests that corporate client decisions are more economically driven than those for personal clients.

4.12. Given that the servicing aspect is something more experienced by the individual, it seems reasonable that a corporate lapse decision will tend to be more financially motivated.

4.13. There were generally more links to the lapse rates in the corporate business than for the other two lines of business, up to and including September 2014. However, after this period there were very few links at all, suggesting that the system had changed and the drivers in our analysis were no longer driving lapse rates.

4.14. The section below titled Tipping points investigates this in more detail.

BETWEEN LAGGED LINES OF BUSINESS

4.15. As mentioned, we decided to lag the data to see if this resulted in any additional drivers being visible in the analysis. The lags that were applied were for both one and three months, in order to capture a range of potential policyholder behaviours. A lag of one month was intended to reflect a policyholder's knee-jerk reaction to an event, where something happens and he or she very quickly decides to lapse the policy in response to that event. In contrast, the three-month lag was intended to reflect a more considered response, where an event occurs to make the policyholder unsatisfied, but he or she carefully considers their options before making a final decision to lapse the policy.

ISA

4.16. When analysing the lagged data for the ISA business, we found that, once again, the drivers were inconclusive. Again, this reinforces the conclusion that policyholders behave much more personally with their ISA investments and tend to withdraw the money when something requires them to, rather than in response to a trend.

UNIT TRUST (PERSONAL)

4.17. For the personal unit trust business, lagging the data reinforced the idea that the service drivers were the key reasons for people lapsing. In addition to the drivers without the lags, concepts such as number of written enquiries, number of complaints and compensation paid were also seen to be drivers.

UNIT TRUST (CORPORATE)

4.18. When the data was lagged on the corporate business, there were not many drivers that showed up on the analysis. This seems to suggest that the corporate business reacts more quickly to changes in circumstance than personal business, which is in line with how the businesses would be expected to react.

4.19. The only exception to this was the addition of some weak links to service-related drivers, which suggested a slightly delayed reaction from the corporate business in response to customer service levels: It presumably takes a little time for staff to complain to an employer about service levels and for the employer to ultimately take action.

DRIVERS OVER TIME

4.20. There were certain instances where it was observed that the drivers of lapse rates changed over time.

4.21. One example of this is the corporate unit trust business. The diagram below shows how the drivers changed over a period of two years:

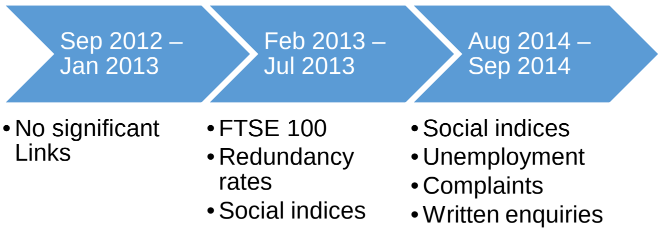

4.22. It would have been interesting to observe how these drivers changed over a longer period to see if the drivers returned to any previous state. It is anticipated that as companies continue to be more conscious of monitoring their data and build up a data bank containing many years' worth of data, this will become more accessible.

TIPPING POINT

4.23. A tipping point event was observed in the analysis and was found in September 2014 for the corporate unit trust business.

4.24. The following diagram highlights the tipping point in red, where the graphical complexity and total uncertainty peak before collapsing to much lower levels (as explained further in sections 3.13-3.15).

FIGURE 4.3: TIPPING POINT IN SEPTEMBER 2014

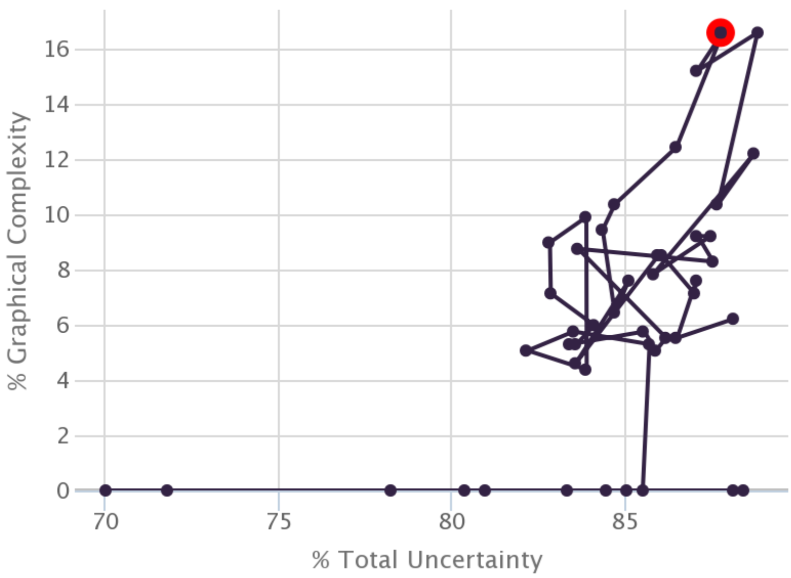

4.25. At this point in time, the system showed high levels of uncertainty (meaning that there were high levels of unpredictability in the system) and high levels of complexity (meaning that the system had high potential to collapse). These two factors combined indicated that the system was at risk of reaching a tipping point.

4.26. The notion of a tipping point in September 2014 for the corporate unit trust business is also reinforced by analysing the number of links to lapse rates, as shown in the diagram below:

FIGURE 4.4: TIPPING POINT IN SEPTEMBER 2014

4.27. A strong peak can be seen in September 2014, and it then disappears the next month.

4.28. One potential reason for this is the fall in global markets in September 2014, often referred to in news outlets as the 'September slump.' Not only would the direct impact be allowed for in the analysis between lapse rates and the market indices analysed, but also the indirect impact on the secondary variables, such as interest rates, unemployment rates and social indices.

WHAT WE DIDN'T FIND

4.29. Interestingly, some of the 'traditional' drivers of lapses were not highlighted by our analysis. This is likely to be because of the less traditional type of business that we were investigating and the unique relationship that the client's advisers have with their clients. However, it was still interesting to observe that result.

4.30. Typically, it might be expected that duration in force and inflation would be key drivers of lapse rates. If this was the case, then we would expect to have seen a distinct pattern between these variables in our analysis. However, as you can see from the image below for the personal unit trust business, this was far from the case.

FIGURE 4.5: ANALYSING TRADITIONAL DRIVERS

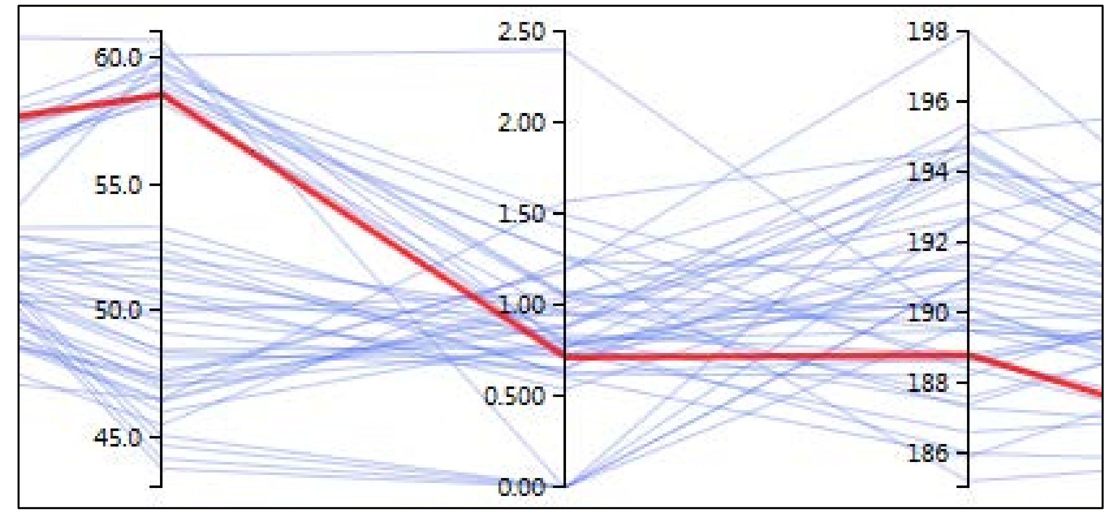

Average duration in force **Lapse rate** **Food RPI**

4.31. From this graph, it is clear that neither lapse rate nor duration have a strong relationship to inflation—the lines linking them are spread across wide ranges of the axis. For lapse and duration, we can see that for durations below 45 months, there are close to zero lapses. Thereafter, the pattern is more complex: Lapses tend to increase for the period up to 48 months, then reduce to a slightly lower level before rising again in just under 60 months. This suggests that different 'regimes' are influencing behaviours, so the relationship with duration is not entirely straightforward.

5. USING THE FINDINGS

5.1. There are many ways in which the client could use the results of the investigations going forward. These include:

- Updating its lapse modelling to reflect its new view on lapse rates based on the results of the research investigation
- Capturing more information about policyholder interactions in a structured way, thereby permitting a more detailed analysis of lapse drivers in future
- Extending the analysis to include additional product types that the client offers
- Prompting the advisers working on the ISA business to develop a better understanding of the policyholder's current personal circumstances, which may help to inform how likely they are to lapse
- Anticipating periods of high lapses based on the current drivers flagged by the analysis and work on customer relationships ahead of these periods
- Amending its product design to be more resistant to the main drivers of lapse rates
- Talking to its advisers about its experiences of policyholders lapsing, and combine this with the mathematical analysis to give a hybrid view on lapses
- Investigating what policyholders did with their lapsed policy to gain additional insight into the policyholder behaviour.

5.2. There are numerous ways in which the analysis could be extended in the future. These include:

- Carrying out investigations into other policyholder behaviours such as customer queries, new business and complaints.
- Lagging the data by other time periods than one or three months. For certain drivers, a longer lag may be more appropriate—for example, a fall in the stock market might cause concern for policyholders, but they might only lapse after six months, when they have seen that the detrimental effect on their policy is not improving. Alternatively, a shorter lag could be considered, such as weekly; however, weekly data would be difficult to obtain for all variables.
- Considering additional drivers in the time series analysis, including drivers which might be more difficult to quantify. Such indicators might be emotive indicators which could include consumer happiness, satisfaction with the product and the public's perception of the company, or could be more market-based such as levels of competition or the market's perception of the company.
- Adding additional companies to the analysis to demonstrate the behaviour in a border range of company types. The client has a somewhat unique distribution model, so analysing a more traditional insurer could add an extra dimension to the analysis.
- Studying alternative product types. It would be interesting to investigate the lapse rates of some non-profit and with-profits products to see if they behave differently to the unit-linked products. Analysing a term assurance product would also add more depth to the analysis, and it would be interesting to compare the behaviour of term assurance customers who would receive no benefit upon lapsing with those of an endowment assurance product who generally receive a surrender value if they lapse.

6. CONCLUSIONS

SUMMARY

6.1. Various key features are confirmed by this research, with plenty of potential to investigate and utilise these conclusions further.

LAPSE BEHAVIOUR VARIES WITH PRODUCT TYPE

6.2. Drivers of lapses vary significantly for the client, depending upon the product in question. The drivers of lapses can be grouped together to inform the sources of data that the client should focus on when analysing its lapse rates:

- For the ISA business, the analysis showed that the drivers of lapses depend very much on the personal circumstances of the policyholder, and so the best way to understand the propensity to lapse may be to develop a better understanding of the policyholder's personal circumstances.
- For the personal unit trust business, the analysis showed that the drivers of lapse rates are based on the service provided and the drivers that have a direct impact on the policyholder themselves, such as unemployment.
- For the corporate business, the analysis showed that the drivers of lapse rates were market-based, suggesting correlation between market and lapse risk.

6.3. The analysis also challenges the traditional drivers used, such as duration in force. Whilst they have proved to be statistically viable predictors for some books of business, our analysis suggests that this may not be due to any actual underlying linkage.

DRIVERS CHANGE OVER TIME

6.4. Drivers can change over time, especially after a tipping point event. Therefore, it is crucial not to expect the groups of drivers seen at one point in time to persist for the long term, but to continue to monitor the system to identify changes in drivers over time.

6.5. Even where there is sufficient past data to validate an assumption statistically, it is worthwhile to understand how the underlying drivers might be changing so that long-term assumptions are set with a full assessment of the dynamics that could be expected to occur over the relevant timeframe. There is also operational benefit to monitoring and understanding how persistency drivers vary throughout the year: Rather than persistency just being an actuarial assumption, it can become a core input to servicing and retention strategy.

6.6. This type of approach is particularly relevant when estimating mass lapse behaviours, where they typically arise due to a breakdown in previous behaviours. Looking at emerging trends to understand when such a tipping-point might be building up could be very useful in creating sufficient time to organise an appropriate intervention or response.

CONTACT

Jennifer Smith
jennifer.smith@milliman.com

Neil Cantle
neil.cantle@milliman.com