DEPARTMENT OF ECONOMICS AND FINANCE

UC BUSINESS SCHOOL

UNIVERSITY OF CANTERBURY

CHRISTCHURCH, NEW ZEALAND

A Lifecycle Approach to Insurance Solvency

Yuechen Dai
Tonghui Xu

WORKING PAPER

No. 13/2021

Department of Economics and Finance
UC Business School
University of Canterbury
Private Bag 4800, Christchurch
New Zealand

WORKING PAPER No. 13/2021

A Lifecycle Approach to Insurance Solvency

Yuechen Dai[1†]
Tonghui Xu[1]

November 2021

Abstract: At present, most well-known insurance regulatory bodies focus on reviewing the solvency of insurance companies within a one-year period. However, the operation of insurance companies is a long-term business, with most policyholders planning on holding a policy over many years, not just one. This research adopts a new perspective for measuring the insolvency risk faced by insurance companies over a longer time period by estimating their full expected lifetime (the number of periods into the future that an insurer can be expected to remain solvent, given their initial capital reserves), which has significance for insurance regulation. This research uses python numerical methods to simulate the operating conditions of insurance companies with different initial reserves, and capture the period in which the company becomes insolvent. The results show that, as is logical, the higher is the initial reserve fund, the longer one can expect the company will be in business before insolvency. In addition, our simulation model helps to explain how the relevant probability density for the insolvency date, given an initial reserve fund, can be estimated. By comparing different probability density functions, we find that a lognormal density form provides a reasonable starting point for the density in question.

Keywords: Insurance regulation, simulation, insolvency

JEL Classifications: C02, C15, C63

[2] Department of Economics and Finance, University of Canterbury, NEW ZEALAND

† Corresponding author: Yuechen Dai. Email: yuechen.dai@pg.canterbury.ac.nz

A Lifecycle Approach to Insurance Solvency

Yuechen Dai and Tonghui Xu

1. Introduction

By its very nature, insurance is a risky business. Aside from the obvious risks derived from the basic business of underwriting the risky situations of policyholders, insurance companies also face risks to their financial stability deriving from their investment activities. Insurance policyholders enter into agreements with their insurers that involve up-front premium payments in exchange for potential indemnity receipts at some future date. While the premium payment is non-risky, there is an inherent risk around the ability of an insurer to always honor their indemnity commitments. It is for this reason that the insurance sector is regulated, with the main issue of concern being financial solvency.

The insurance industry has yet to form a global unified solvency supervision standard. Plantin & Rochet (2009) analyzed the bankruptcy of four insurance companies in the 1990s and made recommendations for monitoring the solvency of insurance companies, including delegating monitoring to a prudential agency. In recent years, some countries and regions with developed insurance industries have taken the initiative to learn from the achievements of the Basel Accords for the banking sector in the process of formulating solvency supervision standards for insurers, since many risks faced by banks and insurance companies are similar.

In the insurance regulatory systems, capital adequacy regulation occupies a core position. Solvency II, currently being implemented in the European Union, fully draws on the "three-pillar" regulatory system of Basel II, which gradually converges the regulatory concepts of insurance and banks, and promotes the consistency of financial supervision. The first pillar focuses on minimum capital requirements (MCR) and solvency capital requirements (SCR). The MCR is the minimum capital required by the regulatory authorities for insurance companies to maintain their normal solvency even in the face of adverse market conditions. It is intended to correspond to an 85% probability of adequacy over a one-year period. Second, SCR is the capital held by an insurance company in order to deal with major unforeseen losses and guarantee compensation to policyholders over the next one-year period with a probability of at least 99.5%. As can be seen, both of these standards work with a horizon of a single year.

The operation of insurance companies is a long-term business, with most policyholders planning on holding a policy over many years, not just one.[1] However, in Solvency II, the risk and thus the capital for the undertaking insurance and reinsurance companies only need to be evaluated on a one-year time horizon (Ferriero, 2016). The same issue also appears in U.S. insurance regulations. The National Association of Insurance Commissioners (NAIC) adopted the Financial Analysis and Surveillance Tracking (FAST) solvency monitoring system and risk-based capital (RBC) requirements. However, a limitation of both of them is that they are static rather than dynamic approaches to solvency testing (Cummins et al. 1999). With this in mind, the present article considers insurance solvency from the perspective of the financial survival of an insurer over its entire lifecycle rather than just a single period into the future.

The reserve system is widely used in the economic and financial industries. Similar to the bank's reserve system, the establishment of an insurance company also requires a reserve fund to be held. In order to guarantee the normal operation of insurance companies and to protect the interests of the insureds, countries have generally stipulated in their legislation that insurance companies should hold a certain amount of liquid reserves to ensure that insurance companies have sufficient solvency for the scale of their insurance business. At present, much research has shown that unreasonable allocation of reserves is one of the reasons leading to the insolvency of insurance companies. Leadbetter & Dibra (2008) conducted a comprehensive study on the reasons for involuntary market exit in the property and casualty insurance industry in Canada from 1960 to 2005. There is evidence that insufficient reserves are one of the important reasons leading to the insolvency of these insurance companies. Carson & Hoyt (2000) identified important variables that cause life insurance companies to have financial distress in EU. The results indicate that high capital reserves and good profit status can reduce the probability of insolvency of insurance companies since they can provide a buffer against larger than expected losses or smaller than expected investment gains. Similarly, in the field of reinsurance, Cai et

[1] For example, life insurance is typically a very long-term arrangement.

al. (2014) provide a reinsurance risk model that incorporates the regulatory requirements on the initial reserve of a reinsurance contract seller and the possible default by the seller. Their results show that the regulatory initial reserve and the default risk have a significant impact on the optimal reinsurance strategy.

The general perception of policyholders is that the expected solvent life of an insurance company increases the larger is the initial reserve (ceteris paribus). This implies that an insurance company with more initial reserves will be more attractive to a potential policyholder of a risk that could play out over many future periods. In order to study whether this view is correct, we establish a simple lifecycle model of an insurance company. An insurance company sells insurance contracts at a given premium at the start of each period of operation, and at the same time bears a random number of indemnity payouts. If the total amount of indemnity payments in a given period is greater than the sum of the initial reserve fund and premiums received in that period, the insurance company will face insolvency. However, when a potential policyholder initiates a contract with an insurer, that policyholder is interested in more than simply the solvency situation for that one period, rather, of importance is the expected solvent lifetime of the insurer. For life insurance, the policy holding period is longer, which requires life insurance companies to have a longer life cycle. If the life cycle of an insurance company is short and cannot survive for a long time to fulfill its promise to the policyholder, then will not only cause losses to the policyholder, but also cause hidden dangers to the insurance market. For insurance companies with long-term policy holdings, it is very necessary to study their life cycle. However, there is little research dedicated to the study of the relationship between the initial capital of an insurance company and its overall lifecycle under dynamic conditions. Our research hopes to fill that gap.

In order to study the relationship between the lifecycle of the insurance company and their initial capital fund, we use a simulation program to perform numerical simulations, since the model involves many complexities that make it difficult to adopt a purely theoretical model. Simulation can provide valuable insights into problems of this sort. At present, simulation programs are widely used in research concerning the insurance industry. Cummins et al. (1999)

analyzed the main models used by U.S. insurance regulators to predict insolvency in the property-liability insurance industry and came up with a relatively new solvency test method – cash flow simulation. Their results indicate that dynamic financial analysis based on cash flow simulations seems promising to provide regulators with better forecasts of insurance companies' solvency. Alm (2015) used a simulation program to build a model that can generate non-life insurance risk solvency capital requirements in order to solve the problem of how to aggregate risks of single insurance types, lines of business or risk classes. Through numerical simulation, we can easily generate the flow trajectory of the insurance company's reserves and record the time point of its insolvency.

Our research adopts a new perspective for measuring the insolvency risk faced by insurance companies over a longer time period by estimating their full expected lifetime (the number of periods into the future that an insurer can be expected to remain solvent, given their initial capital reserves), which has significance for insurance regulation. Our focus, then, is on finding the probability density for the first date of insolvency of an insurer, as a function of the initial reserve fund (i.e. the reserve fund at date 0, which is the date at which we assume a policyholder is considering initiating an insurance arrangement).

The remainder of the essay is organized as follows. Section 2 provides the theoretical framework and discusses the main assumptions behind the simulation that we use. Section 3 provides the results of the simulations that we ran, including the estimated lifecycle densities. Finally, section 4 concludes and discusses opportunities for future research.

2. Theoretical Framework

2.1 Assumptions

We assume that the solvent life of our subject insurance company is t periods, which is a variable (rather than a fixed parameter). In each period of operation, the insurer has n identical policyholders, each of which suffers a loss of value L with probability p. We assume a perfectly competitive insurance market in which all policyholders have the same contract, namely full coverage at a fair premium pL. Thus, the expected profit of the insurance company is 0, conditional upon being solvent.[2]

The insurance company has initial reserve funds in period 0 of $R_0 > 0$. Over time, the reserve fund value fluctuates along a random walk according to the amount of indemnities paid. In any period in which total premium income is greater than total indemnities, the profit is added to the reserve funds. And if the total premium income is less than total indemnities, the loss is financed out of the reserve funds. In the first period in which total indemnity claims exceed premium income plus the reserve funds, the company defaults and pays each indemnity claimant an equal share of current premium income plus reserve funds.[3] Then the insurance company closes down due to finding itself insolvent.

In order to focus on the problem at hand, we assume that over time the insurer neither adds to the reserve fund (outside of any accumulated profit), nor subtracts from it (e.g. for payments of dividends). Incorporating each of those would affect the final results in very obvious ways, and in our conclusions we suggest that some attention to these possibilities might be fruitful as future research. We also assume that the reserve fund is not invested in any interest-bearing asset. This avoids the complications of additional risks, and again, including such options affects the outcome of the simulation process in a very obvious and predictable way.[4]

2.2 Insolvency probability in one period

In the model, in each time period, premium income is $q = npL$, and the total amount of

---

[2] This assumption is purely for convenience. The model can equally accommodate an assumption of a full coverage contract that is sold at a higher than fair premium (but at a premium that is no greater than expected indemnity plus the policyholder's risk premium). The effect can be easily seen to simply prolong the life expectancy of the insurer. This effect is, essentially, identical to simply assuming a higher initial reserve fund.

[3] Of course, this means that should insolvency happen, the insurance contract is no longer "fair" for the policyholders since the indemnity will fall short of the amount contracted. It is only priced fairly if the insurer is solvent. There is a significant literature that looks at the effect of the possibility of insolvency on insurance demand. We ignore this effect here simply because it will be a very small, bordering on insignificant, effect when the insurer has a long enough expected life, and when insolvency could imply indemnity payments that are very close to what is owed. In short, full coverage at an actuarially fair premium in each solvent period provides the policyholder with far greater utility than not insuring, even with the potential for an indemnity that is less than the amount expected.

[4] Essentially, our model can be thought of as a baseline, upon which other assumptions can be added.

indemnity claims is a random variable given by a standard binomial distribution with probability p, trial value L, and n trails.

The probability that there will be j indemnity claims in any given period is

$$z_j = \frac{n!}{j!\,(n-j)!} p^j (1-p)^{n-j} \text{ for } j = 0,1,2 \ldots, n.$$

Each claim has value L, so the expected claims cost in any period is

$$EC_i = \sum_{j=0}^{n} L \times j \times z_j = L \sum_{j=0}^{n} j \times z_j$$

But it is well-known that, for the case of standard binomial distribution, this is just $EC_i = L \times n \times p$. Therefore, in each period of solvency, expected profit (the premium income less expected claims) is exactly 0.

This model then defines a random trajectory, which is captured by the size of the reserve fund overtime. Insolvency happens in the first period in which the size of the reserve fund takes a negative value. We can model this process under a simulation, to see how changing the variables, which are p, n, L, and R_0, affects the trajectory. In our simulations, we are only concerned with changes in the value of the initial reserve fund, R_0.

The cumulative distribution corresponding to the density of indemnity claims is

$$G(k) = prob(j \leq k) = \sum_{j=0}^{k} \frac{n!}{j!\,(n-j)!} p^j (1-p)^{n-j} \text{ for } k = 0,1,2 \ldots, n$$

Here, $G(k)$ tells us the probability that there are no more than k indemnity claims in a given time period. Therefore, the probability that there will be more than k claims in any period is just $1 - G(k)$. Insolvency happens in period i if there are at least k claims such that

$$kL > npL + R_i \rightarrow k > np + \frac{R_i}{L}$$

Given this, set $k_i^* = np + \frac{R_i}{L}$, so that in any period i, the number of claims that triggers insolvency is k_i^*. In any period in which the reserve fund satisfies $R_i < (1-p)nL$, we know that

$$k_i^* = np + \frac{R_i}{L} < np + \frac{(1-p)nL}{L} = np + (1-p)n = n$$

So, in all such periods, $k_i^* < n$, there is some positive chance of insolvency.

We then know that the probability of insolvency is

$$s_i = 1 - \sum_{j=0}^{k_i^*} \frac{n!}{j!\,(n-j)!} p^j (1-p)^{n-j}$$

And the probability of solvency is $1 - s_i$.

So, in this model, in any period of solvency, the expected profit is 0. And by definition, in any period of insolvency, the expected profit is also 0. Therefore, the total expected profit of this insurer is 0.

3. Simulation model

In our program, we define n to be the number of policyholders in every time period, p is the probability of the insured accident happening in each period, i is period counter, L is the value of the potential loss, and R is amount of initial reserve funds. We assume that n, p and L are all constant over time. We then generate a binomial discrete random variable for the number of claims that occur in each period i. When the total amount of indemnity claims exceeds the premium income plus the reserve funds, the insurance company becomes insolvent. Using randomly generated data of the number of claims, we can calculate the probability of insolvency in each time period i, and more interestingly, the density for the lifecycle of the insurance company.

For all of our simulations, we set $n = 1000$, $p = 0.1$, $L = 500$ and $i = 1000$. We run three simulations, with different initial reserve fund values. Concretely, we assume reserve fund values of \$40,000, \$80,000 and \$120,000. For each reserve fund value, the simulation calculates the insolvency period 100,000 times, in order to obtain an estimate of the probability distribution of the lifecycle of the insurance company.

3.1 Some illustrative cases

We let each simulation run for a maximum of 1000 “periods”. The date at which the reserve fund goes below 0 is when the simulation stops (unless that does not happen before 1000 periods, in which case the simulation simply stops at period 1000), and the insolvency date is recorded. The following graphs show a few illustrative examples (i.e. selected single runs out of the 100,000 that we undertake) for each of the three cases studied. Each following figure is divided into two parts; the upper part is the evolution of insurance company reserves over time and the lower part is its one-period insolvency probability over time.

Cases with initial reserve fund of $40k:

Case 1: This is an example in which the company had about 30 periods of losses around expected value, then about 25 periods in which losses on average exceeded expected value. Around period 60, the probability of insolvency jumps to about 0.1, at the same time as the accumulated reserve fund drops to its lowest level of around $10k. But then the company recovers (losses below expected value), the reserve fund recovers up to about $30k, but then again drops. At about period 80, the probability of insolvency again starts to rise above 0, and it reaches about 0.4 at about period 90 (reserve fund now below 10k), before another period of recovery. But after about period 105, the company suffers a gradual demise, until at period 119, a major loss event puts the company into insolvency.

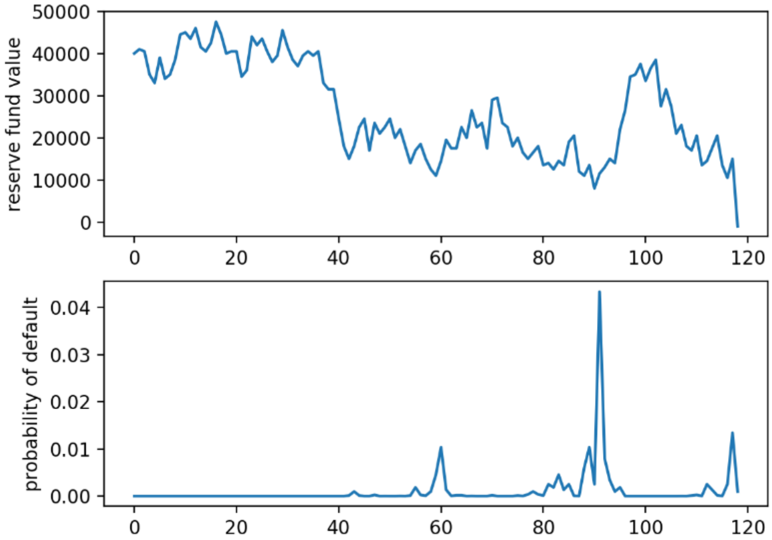

Figure 1: Case 1 of the insurer with initial reserves of \$40k

Case 2: This case shows a company that was generally rather unstable right from the start. Only by about the fourth period of operation, the accumulated reserve fund had dropped from \$40k to only half of that amount. Aside from the recovery periods around periods 15 to 17, this company was generally rather unlucky. Like the first case study, we see that the probability of insolvency only starts to show positive when the accumulated reserve fund goes as low as \$10k. The company managed to scrape through on a tiny reserve fund for about 4 or 5 periods at the end, before hitting its bankruptcy event in period 30. Never-the-less, again the probability of insolvency for the period in which insolvency occurred was only about 0.2, and certainly was not at an all-time high for this company.

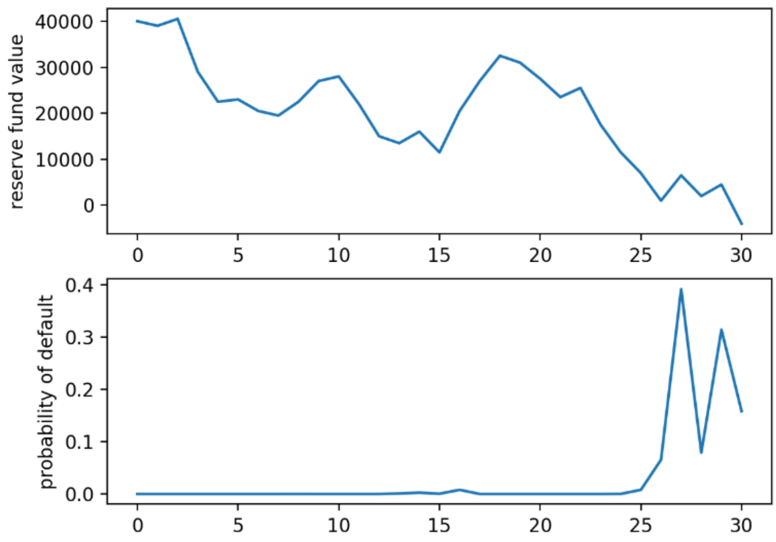

Figure 2: Case 2 of the insurer with initial reserves of $40k

Case 3: This is the case of a fortunate company, that managed to remain solvent for a very long time (about 375 periods) on an initial reserve fund of only $40k.[5] The company suffered a minor scare at around period 40 (where once again the reserve fund dropped to around $10k), but then had a very long period of about 250 periods in which indemnity payouts generally were less than premium income, and the accumulated reserve fund over this period went even well above $100k. The demise began at around period 275, with a recovery at about period 300, before the eventual collapse. Because the accumulated reserve fund was so high just after period 300, the collapse actually took quite a few periods to play out, but in the end insolvency was reached after about 10 periods of significant losses before closure at about period 375. Once again, looking only at the periodic probability of insolvency, there were no red-flags for this company right up to the very insolvency date (the probability of insolvency was indistinguishable from 0 right before insolvency occurred).

[5] This was not the company that lasted the longest on a $40k initial fund. In another example, a $40k initial fund company was solvent for close to 700 periods!

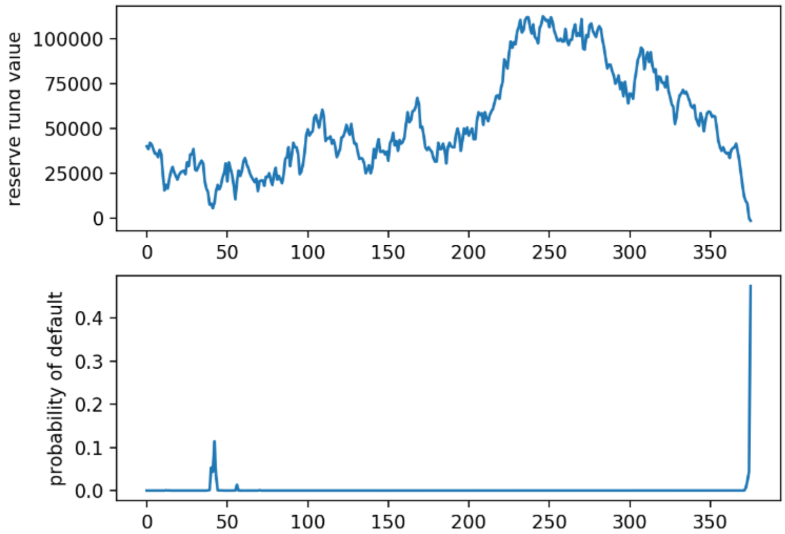

Figure 3: Case 3 of the insurer with initial reserves of $40k

Cases with initial reserve fund of $80k

Case 1: With an initial fund of $80k, it is much more likely that the company will survive for quite some number of periods. This first case shows a typical evolution of the accumulated reserve fund and the probability of insolvency corresponding to an initial reserve fund of $80k. As can be seen, the company experienced very minor average losses for 150 periods, at which time the accumulated reserve fund had dropped only by about $20k from its initial value. What has killed this particular company is the 25 periods of bad luck between periods 150 and 175, over which time the accumulated reserve fund went from about $60k to about $10k. As is the case for most of our simulations, having a reserve fund of about $10k or less causes a spike in probability of insolvency to strictly positive levels, and this can be seen for the case of this particular company, when probability of insolvency shoots up to about 0.25 as soon as the reserve fund hits numbers below $10k. There is a short period (maybe 10 simulation periods) of recovery, followed then by the collapse with about 10 periods of continual indemnity payouts exceeding premium income, until insolvency occurs at about period 190.

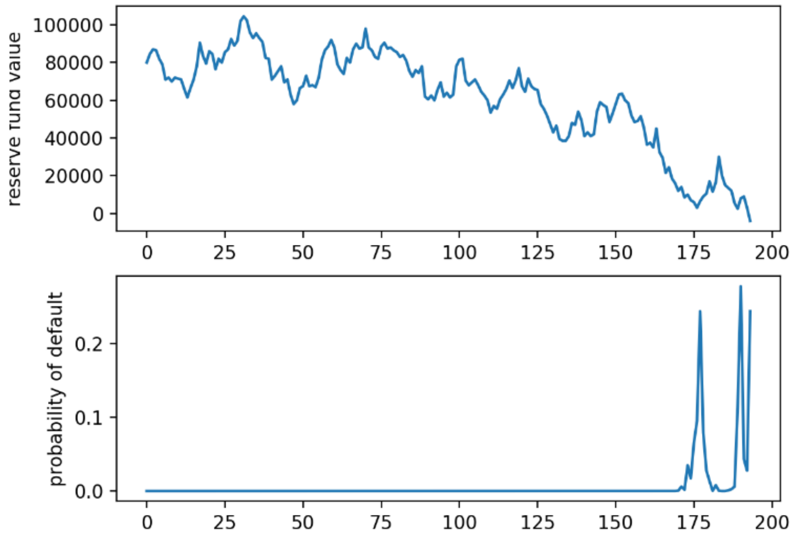

Figure 4: Case 1 of the insurer with initial reserves of $80k

Case 2: In this example we have again a relatively fortunate insurer with an $80k initial fund that managed to last around 170 periods before succumbing to insolvency. The case shows a typical evolution, where upon the company experiences, first, a long time of prosperity (approximately 70 periods) in which the reserve fund increases up to about $120k. This is then followed by about 50 periods of gradual recession that wipes around $100k from the reserve fund. There is then about 10 periods of recovery, in which the reserve fund almost doubles from its low point. Finally, over about 30 periods, the firm succumbs to a final insolvency at around period 170. Of note is that over this entire lifecycle, the probability of insolvency in the next period only went positive right before insolvency. This is a common feature through many of our simulations, where this probability only goes notably positive when the reserve fund reaches about $10k, from which point it often only takes a more few periods for insolvency to occur.

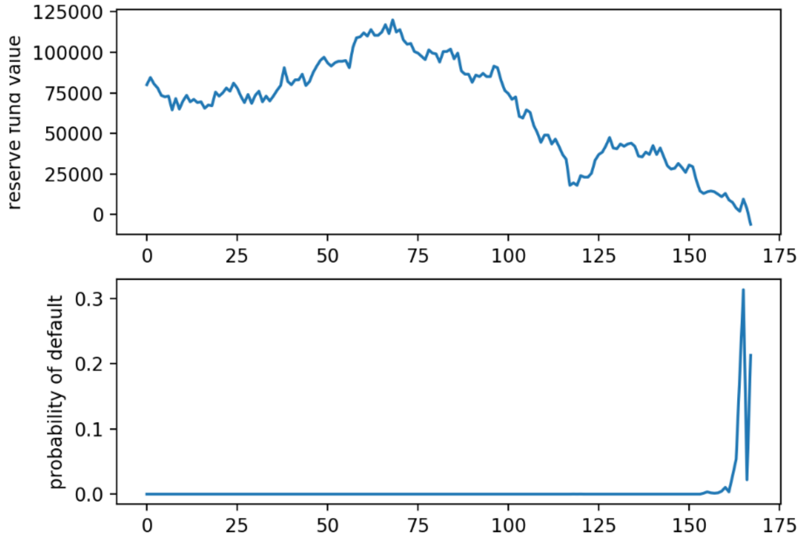

Figure 5: Case 2 of the insurer with initial reserves of $80k

Cases with initial reserve fund of $120k:

Given the high initial reserve fund value, most of these companies experienced a significantly long total life-time, but still all of them went insolvent in the end. A typical evolution of the reserve fund showed many periods of very gradual average demise, but with no discernable probability of insolvency, until a sequence of catastrophic indemnity periods causes insolvency over a period of only about 10 or 20 simulation periods.

<u>Case 1</u>: Here we have the case of an insurer for whom the reserve fund undergoes cycles of expansion and recession for about 210 consecutive simulation periods. But after that, there is a gradual but continual demise, and over the next 120 or so simulation periods, the accumulated reserve fund drops from about $120k to about $30k. Nevertheless, this recessionary period does not show at all in the probability of insolvency, which remains at essentially 0. Looking only at the periodic probability of insolvency, there would be no indication that this company was in any danger until immediately before it actually becomes insolvent. The collapse occurs over about only 5 simulation periods, in which indemnities continually exceeded premium income, causing insolvency at around period 345.

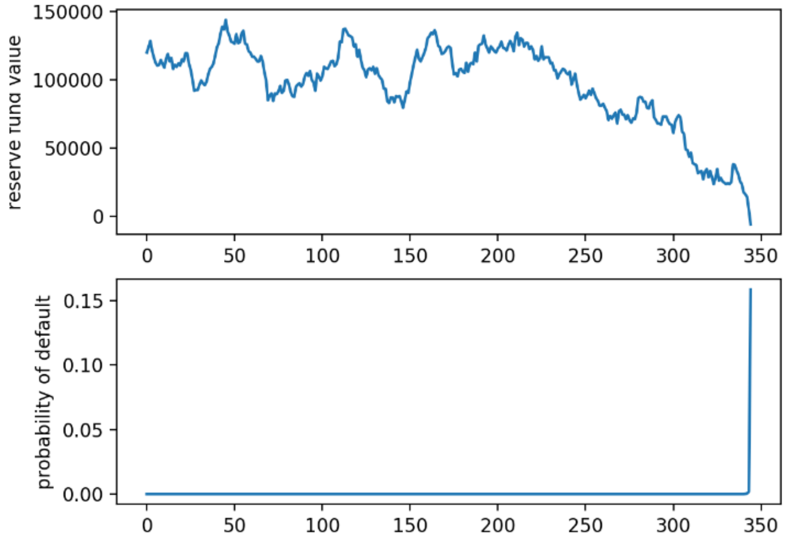

Figure 6: Case 1 of the insurer with initial reserves of $120k

Case 2: In this case we have an insurer that began with a series of unfortunate periods, losing about $45k of its initial reserve fund over only about 15 periods. There was a short period of recovery, where the accumulated reserve fund went close to $100k, but then a long period of continual losses on average. Between about simulation periods 30 and 80, this company had a long series of unfortunate outcomes, and the accumulated reserve fund fell from around $100k to under $10k, causing the probability of insolvency to spike up to numbers between about 0.05 and 0.1 for about 10 consecutive simulation periods. However, the company survived, and indeed prospered for close to another 100 simulation periods, before succumbing to a gradual demise and eventual insolvency.

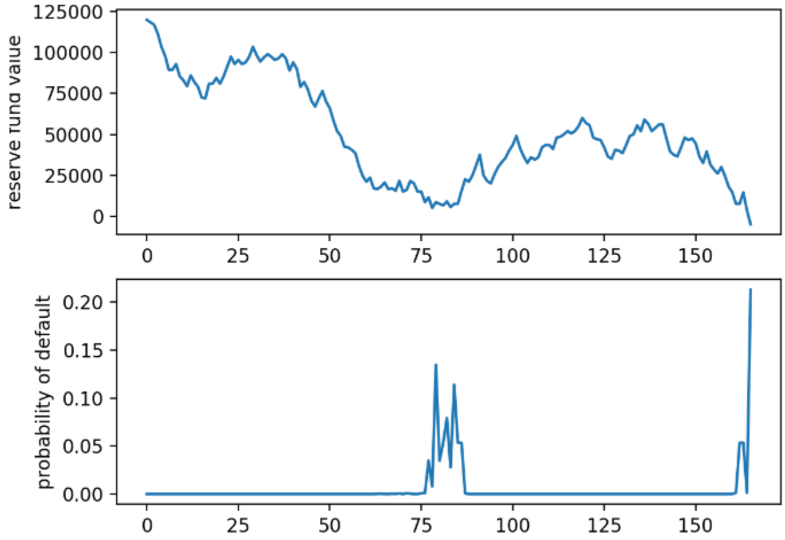

Figure 7: Case 2 of the insurer with initial reserves of $120k

Case 3: In the final case example that we present for an initial fund of $120k, we have a company that enjoyed a very long history of solvency, close to the maximum of 1000 simulation periods. In its heyday, at around period 480, the company had accumulated close to $240k in reserve fund, essentially doubling the initial investment. With such a high reserve fund, it took this company almost 500 more simulation periods before its eventual collapse into insolvency. As can be seen in the lower panel, for all but the very last few periods, this company had no discernable probability of insolvency in any period.

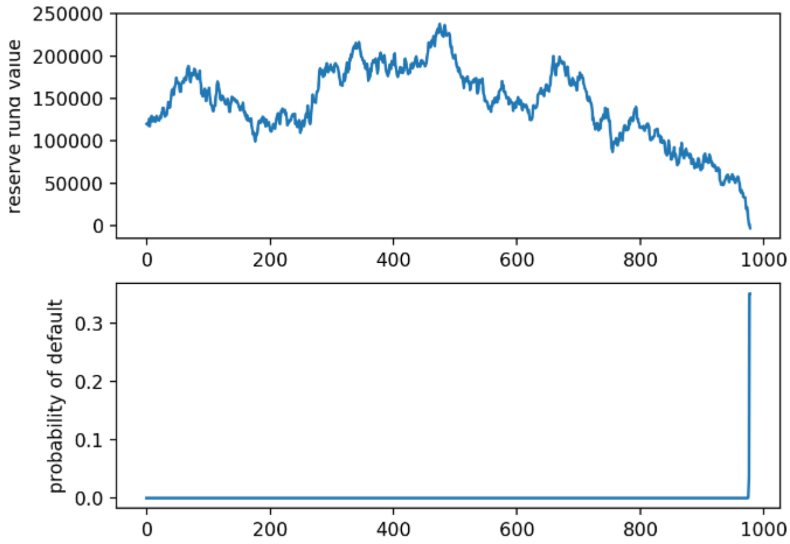

Figure 8: Case 3 of the insurer with initial reserves of $120k

In order to verify whether more initial reserves can extend the life expectancy of the insurance company, we generated the life expectancy of the insurance company under three different initial reserve conditions through the simulation results. The expected values of the three simulations are displayed in the following table:

Initial fund size	Expected period of insolvency
$40,000	202.18
$80,000	347.12
$120,000	462.78

Table 1: The expected period of insolvency of the three simulations

This is shown in the following graph, along with the 2^{nd} order polynomial least squares regression line through the data:

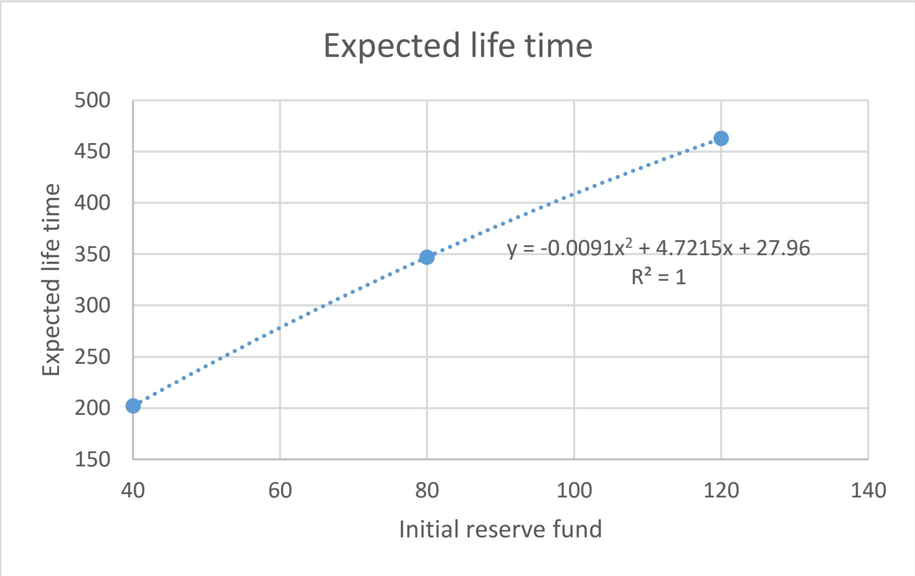

Figure 9: The expected life time of insurer as a function of initial reserve fund

This graph indicates that the expected insolvency date, given an initial reserve fund value of F, is given by $EL = -0.0091F^2 + 4.7215F + 27.96$. It is a very slightly concave function, indicating that there appears to be slightly decreasing returns (in terms of expected life time) to increases in the initial fund value. However, this picture is only drawn based on the results of three different initial reserves. In future research, we will conduct more simulations to use more different initial reserves to refine this figure to verify our conclusions.

The above cases are examples of what can happen, and the actual numbers involved are only relevant in as much as how they compare with each other. The few notable lessons to learn from these case studies are

1. The higher is the initial fund, the longer it can be expected to take before insolvency occurs.
2. Regardless of the size of the initial fund, insolvency is very likely to happen at some point. In some periods, although the one-period insolvency probability is relatively high, the insurance company still survives. However, in other periods, the insurance company goes bankrupt when the one-period insolvency probability is relatively low. Insolvency sometimes was caused by a single huge loss event, and would not have been predicted by the one-period insolvency probability right before the collapse.

3. With our parameter values on the number of policyholders, the size of each loss, and the probability of loss, it happens that the probability of insolvency only becomes discernibly positive when the reserve fund hits numbers around $10k.

3.2 Probability densities for solvency life

The above case studies point to the very logical result that the higher is the initial reserve fund, the longer one can expect the company will be in business before insolvency, and one-period insolvency probability doesn't make good predictions about the collapse that the insurance company is about to face. Our simulations are able to help to show a good estimate of the probability density for the year of insolvency for different choices of initial reserve fund which provides a new idea for predicting the solvency of insurance companies. The following graphs show our three simulations, with reserve funds of $40k, $80k and $120k respectively. Each simulation still allows for a maximum life of 1000 simulation periods, and each simulation was run 100,000 times (that is, for each different initial reserve fund value, we have 100,000 observations on the insolvency date of the insurer). Each one of those simulation trials generates an insolvency year, and the scatter plots show the full set of end-dates on these 100,000 simulation trials, with the relative frequency of each date on the vertical axis. The resulting sets of points are clearly something that we would expect from independent drawings from three different given probability density functions.

From the sets of points, we can get a probability distribution of the life-cycle of the insurance company and it is obviously not symmetric. Our objective is to find a probability density function that can fit this data well. All values of the function must be positive (because probability cannot be negative). Moreover, we can also observe the possible shape of the distribution: initially, the probability presents a relatively rapid upward trend, and reaches a peak, and then presents a slow downward trend, and with a long right tail. This distribution is very similar to the log-normal distribution. Therefore, in order to approximate the underlying density, we have assumed it to be a log-normal density, and we calculated the best-fit density under maximum likelihood estimation (MLE) to the set of simulation outcome points.[6] Those

---

[6] We also tried fitting a gamma density to the data, but the fit was notably worse than the log-normal case.

best-fit log-normal densities are shown along with the full data set from the three simulations.

We also fitted an order 50 polynomial to the data under least-squares. However, of course an order 50 polynomial is of very little practical use, although it does provide a very close estimate for the data. The reason is that its formula is too complicated, and there is a negative probability that does not match the reality. We have used the polynomial fit to compare it with the MLE log-normal curves, in order to see that the MLE log-normal form, which is a useful and quite standard density, does indeed provide a relatively accurate description of the data.

To start with, the case of an initial fund of $40,000 is shown in Figures 9 and 10 below. Figure a shows the relative frequency data points, together with the MLE log-normal curve. Figure b shows the MLE log-normal fitted curve together with the order 50 polynomial best-fit curve. As can be seen, the log-normal density fits the data well, except for the points around the peak of the data. Even if the order 50 polynomial density can better fit the point set at the peak, it has a negative probability that it does not match reality. As we will see, this is (at least partially) an issue with the relatively small size of the initial reserve fund. When the initial fund is larger, the MLE log-normal fit around the data peak improves.

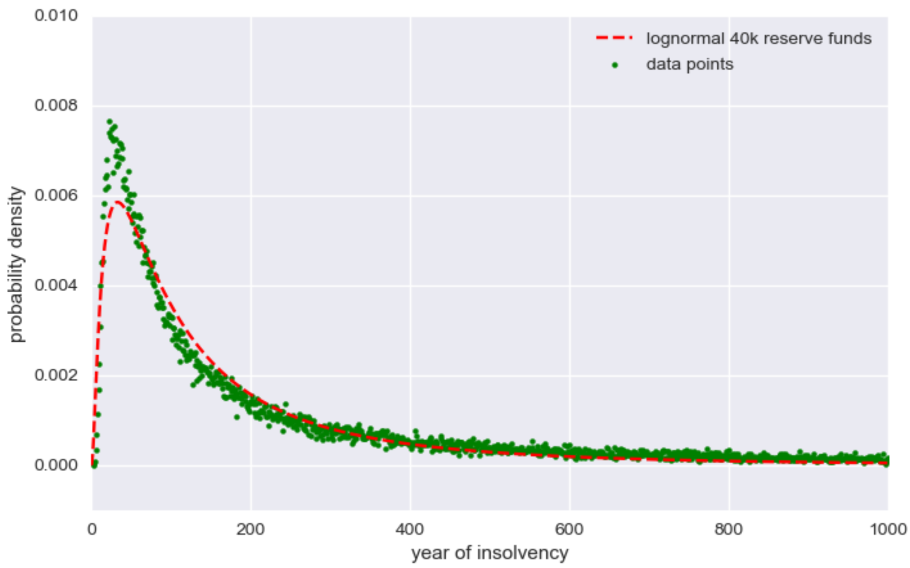

Figure 9: The relative frequency data points together with the MLE log-normal curve

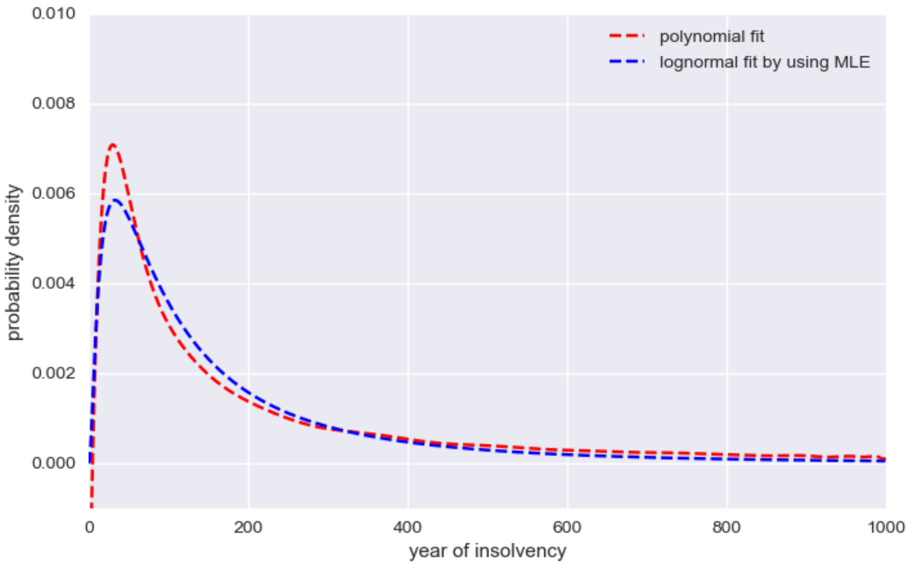

Figure 10: The MLE log-normal fitted curve together with the order 50 polynomial best-fit curve

Figures 11 and 12 show the case of an initial fund of $80,000. What we can see immediately is that the data is more spread out, since the peak density value is lower than with the smaller initial fund. It is also the case that the log-normal density provides a much better overall fit to the data points. Indeed, the log-normal fit tracks almost exactly the order 50 polynomial fit (and of course it avoids the negatively sloped section of the polynomial fit for very small insolvency dates). The order 50 polynomial density does not fit the data very well in the very small insolvency dates part. In figure 12, the order 50 polynomial density showed a downward trend inconsistent with the reality from the beginning.

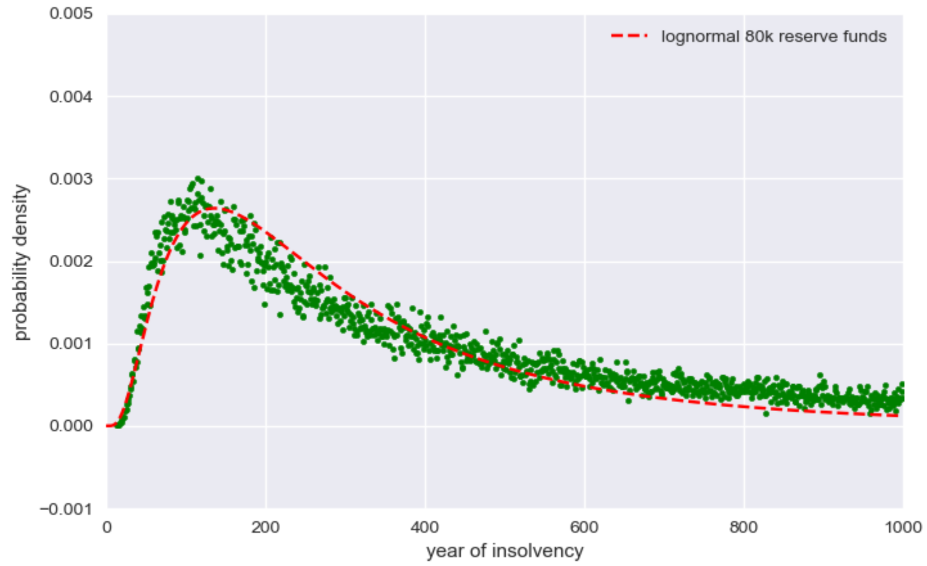

Figure 11: The relative frequency data points together with the MLE log-normal curve

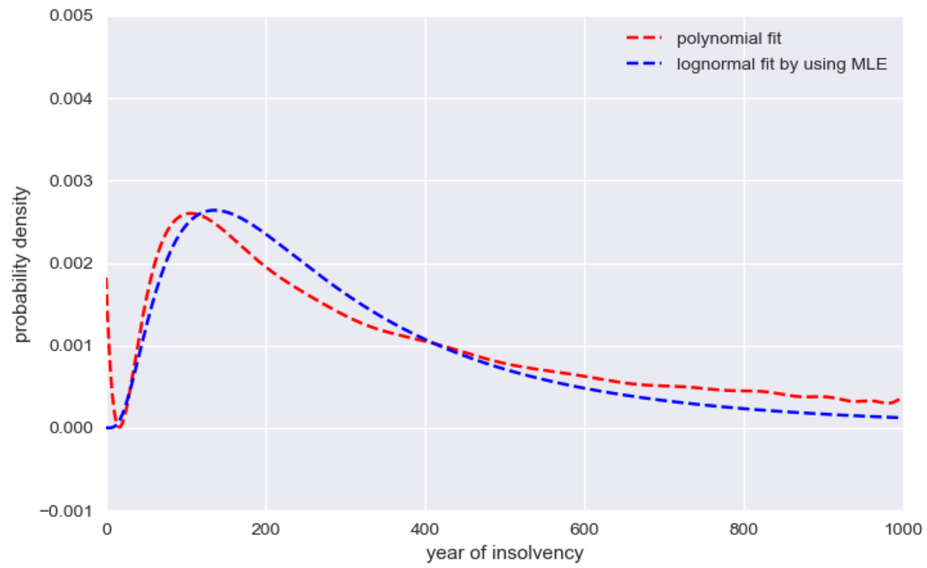

Figure 12: The MLE log-normal fitted curve together with the order 50 polynomial best-fit curve

Finally, in Figures 13 and 14 we have the case of an initial reserve fund of $120,000. The log-normal fit is now much better aligned with the order 50 polynomial fit (notice that the vertical scale on the graph has been reduced in order to better see the two graphs together). Similar to the situation in the first two cases, the 50 order polynomial density cannot fit the data well in the very small insolvency dates since it has a negatively sloped section.

Figure 13: The relative frequency data points together with the MLE log-normal curve

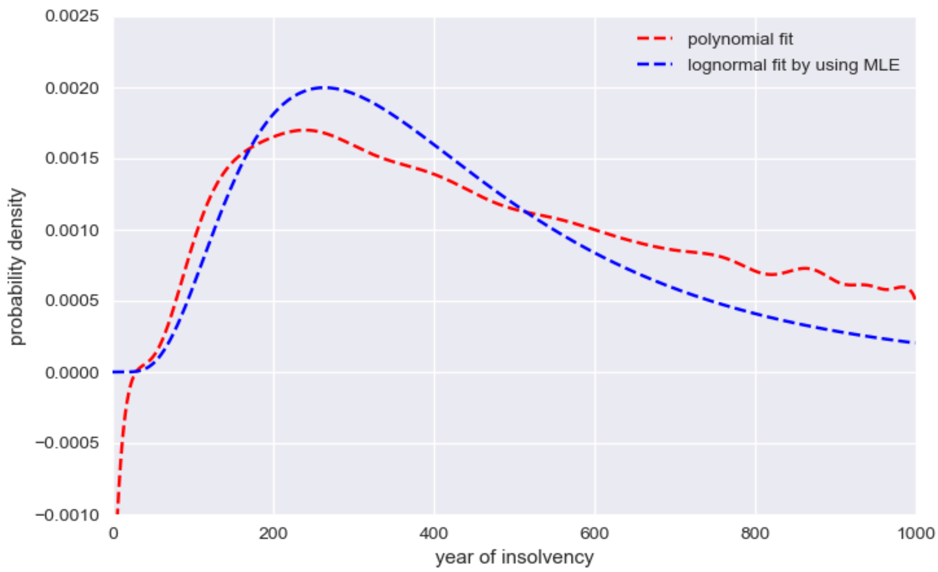

Figure 14: The MLE log-normal fitted curve together with the order 50 polynomial best-fit curve

Figure 15 shows the three log-normal densities together. As can be seen, increasing the initial reserve fund has the effect of spreading the data out, so the peak height is lower and the tail is somewhat fatter, and moving the peak height to the left. Essentially, increasing the initial reserve fund has the effect of a first-order stochastic dominant shift in the density. Therefore, increasing the initial reserve will extend the life expectancy of the insurance company to a certain extent, which is in line with our hypothesis.

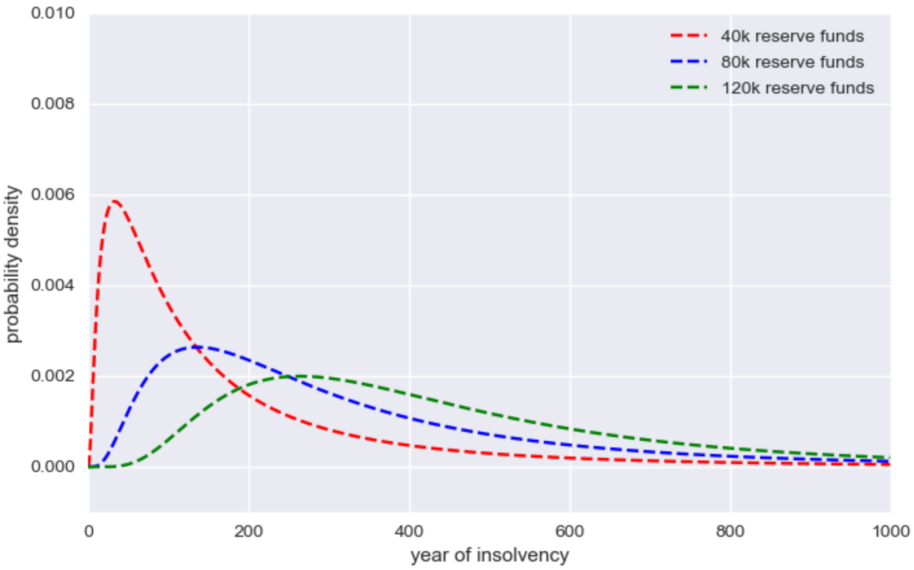

Figure 15: The three log-normal densities curves

4. Conclusions

This research examines the impact of the initial reserve funds of an insurance company on its life cycle. The simulation results show that for insurance companies, more initial reserve funds can provide a higher life expectancy. Indeed, an increase in the initial reserve fund appears to generate a first-order stochastic improvement in the underlying density corresponding to the insolvency date for the insurer. Our model and results provide a new perspective on how insurance insolvency can, and should, be handled by regulators. Above all, our results point to a regulatory standard that is based upon a single-period probability of insolvency tells a very partial story about insurance solvency in general. Rather, we think that what is of the essence is a perspective involving a longer time horizon, with the objective being the probability density for the insolvency date and how that is affected by the insurer's available reserve funds.

Our simulation model helps to explain how the relevant probability density for the insolvency date, given an initial reserve fund, can be estimated. We find that a lognormal density form provides a reasonable starting point for the density in question. While we only ran three concrete simulations (with three different reserve fund values), it would be a relatively simple (even if time-consuming) task to run many more simulations, with many more initial reserve funds, in order to better estimate the curve showing the effect of reserve fund size upon the expected life time of the insurer. This is work that we leave for now on the research agenda.

Our model is based on several simplifying assumptions, any one of which might be relaxed in potential future work on this topic. Of course, the reality of insurance is far more complex than what we have assumed, so our model and results are only a starting point. Extensions to our model could include allowing the insurer the option of adding to or subtracting from the reserve fund over time (besides the natural increases or decreases in the fund due to normal insurance operations). Second, it could easily be assumed that the reserve fund is held in some form of relatively liquid asset, but one that bears some interest. That extension would make no major difference to our results if the investment was risk-free (indeed, allowing the reserve fund to generate risk-free interest is essentially the same as starting the model with a larger reserve fund). Similarly, it would certainly be straight-forward to assume that the insurance product that each policyholder has actually generates some positive expected profit. Again, this really is not particularly different from an assumption of a larger initial reserve fund. We have preferred to simply work on a baseline model, showing how the density that we are interested in can be estimated. Adding in these additional complexities will require further ad-hoc assumptions around the amount of interest that the fund can generate, and the loading on the insurance product if it is to generate positive expected profits.

Another, more fruitful, avenue for future research is to add into the model the possibility that the insurer can re-insure its indemnity claims. Of course, re-insurance could quite possibly be structured such that the primary insurer is solvent always, but that would depend upon the conditions under which reinsurance is written.

Finally, while we tried several different density function forms before settling on the lognormal, it is relatively clear that our lognormal assumption does still miss the mark slightly as an estimator of the true underlying density. Therefore, future work could be centered upon refining the functional form for the density.

References:

Alm, J. (2015). A simulation model for calculating solvency capital requirements for non-life insurance risk. Scandinavian Actuarial Journal, 2015(2), 107-123.

Cai, J., Lemieux, C., & Liu, F. (2014). Optimal reinsurance with regulatory initial capital and default risk. Insurance: Mathematics and Economics, 57, 13-24.

Carson, J., & Hoyt, R. (2000). Evaluating the risk of life insurer insolvency: implications from the US for the European Union. *Journal of Multinational Financial Management*, *10*(3-4), 297-314.

Cummins, J. D., Grace, M. F., & Phillips, R. D. (1999). Regulatory solvency prediction in property-liability insurance: Risk-based capital, audit ratios, and cash flow simulation. Journal of Risk and Insurance, 417-458.

Ferriero, A. (2016). Solvency capital estimation, reserving cycle and ultimate risk. Insurance: Mathematics and Economics, 68, 162-168.

Leadbetter, D., & Dibra, S. (2008). Why insurers fail: The dynamics of property and casualty insurance insolvency in Canada. The Geneva Papers on Risk and Insurance-Issues and Practice, 33(3), 464-488.

Plantin, G., & Rochet, J. C. (2016). When insurers go bust: an economic analysis of the role and design of prudential regulation. Princeton University Press.

Appendix: The lognormal density parameter estimates

The equation of the lognormal density is

$$L(x;\mu,\sigma) = \frac{1}{x\sigma\sqrt{2\pi}} exp\left(-\frac{(Ln(x)-\mu)^2}{2\sigma^2}\right)$$

Where μ is the mean, and σ is the standard deviation.

For our simulations, the following parameter values were estimated (through maximum likelihood estimation), estimates rounded to 5 significant digits:

	$R_0 = 40{,}000$	$R_0 = 80{,}000$	$R_0 = 120{,}000$
μ	4.73369	5.55869	5.96631
σ	1.11489	0.80856	0.62335