Parameterising capital modelling volatility: allowing for changes in volume

30 August 2024

Neil Gedalla FIA
Principal

+44 (0)20 7432 7780
neil.gedalla@lcp.uk.com

Adam Smylie
Associate Consultant

+44 (0)20 7432 6628
adam.smylie@lcp.uk.com

Jade Lagrue
Associate Consultant

+44 (0)20 7432 6720
jade.lagrue@lcp.uk.com

1. Abstract

A common problem in capital modelling involves dealing with volatility assumptions that are inconsistent with the exposure assumptions from which they were originally derived, eg when rolling parameters forward from year to year, due to timing constraints in the capital modelling process, or when performing sensitivity or scenario testing. In order to obtain meaningful results from modelled output, it is important to use parameters that reflect the business being modelled as accurately as possible. The purpose of this paper is to outline a simple formula that capital modellers can use to adjust volatility parameters based on changes in volume.

We show that, based on historical data, the relationship between reserve volume and volatility of reserve movements is well described by a power law formula of the form $CoV = av^{-b}$, where CoV is the coefficient of variation and v is reserve volume. b is a constant to be found, and a is a constant but is not needed to apply the adjustment formula. The data suggest that suitable values for b fall in the range 0.12 - 0.31, and we suggest 0.22 as a suitable value for a typical use case.

This leads to the finding that insurance risk CoVs can be adjusted for the impact of changing volume using the following formula:

$$CoV_{after} = CoV_{before} * \left(\frac{v_{before}}{v_{after}}\right)^{0.22}$$

In addition, we note that benchmark survey data collected indicates that capital modelling practitioners typically make less allowance for the sensitivity of volatility to volume than is indicated by historical data. Volumes larger than around USD $50m may therefore have CoVs that are overstated, and volumes smaller than this may have CoVs that are understated.

Contents

2. Introduction

2.1. Purpose

The purpose of this paper is to provide a solution to a series of related problems in capital modelling, in which volatility assumptions are inconsistent with the exposure assumptions from which they were originally derived.

Some examples of where this arises are as follows:

- Reserve risk coefficients of variation (CoVs) may be calculated based on experience at the previous year end, but need to be adjusted for future projected reserve volumes.
- Underwriting risk parameters may be calculated based on historical data, with volumes varying over the timeframe considered, but planned future premiums may be outside the range observed historically due to changes in business plans.
- Sensitivity testing may consider changes in volumes and require a systematic way of estimating resulting impact on volatility parameters.
- Volatility parameters may be rolled forward unadjusted from quarter to quarter, or from year to year, if a full re-parameterisation is deemed unnecessary (eg if the risk profile is unchanged) or cannot be performed due to time constraints, but may need adjustments for changes in volume.
- Volatility parameters may be required for a multi-year capital projection, or for scenarios involving changes to planned business volumes.
- Distributions for new (or small) classes of business may be parameterised with reference to an existing (or larger) class.

It will often be impractical to align the parameterisation fully to the exposures modelled. In particular, for the last three examples above, there will typically be little information available on the nature of any new business. However, in order to obtain meaningful results from modelled output, it is important to use parameters that reflect the business being modelled as accurately as possible.

The main objective of this paper is therefore to outline a simple algorithm that capital modellers can use to adjust volatility parameters based on changes in volume.

The previous literature on this specific topic is relatively sparse. We are aware of various attempts to examine the relationship between volume and volatility previously – for example, the short paper *Back to Basics, a simple model of insurance volatility*, published by Paul Hewson on LinkedIn[1] in March 2024 outlines a theoretical approach. We are not aware of any published works that attempt to provide a practical solution to the problem outlined above.

2.2. Outline

We have investigated the relationship between the volatility and volume of insurance losses. The model we have used to investigate the relationship between these two quantities is discussed in section 3.

The bulk of our analysis was performed on the Schedule P dataset from the National Association of Insurance Commissioners (NAIC), which includes claims information for major personal and commercial lines for all property and casualty insurers that write business in the US. This dataset was used because it is large enough to provide credible results, and easy to analyse as the data are provided in a consistent format. A brief description of the data used, along with descriptions of the other datasets analysed, can be found in section 4.

[1] https://www.linkedin.com/posts/paul-hewson-790702b4_back-to-basics-a-simple-model-of-insurance-activity-7170767867783680002-qbzL?utm_source=share&utm_medium=member_desktop

We have performed two analyses on the Schedule P dataset:

- **Bucket analysis** of the relationship between volatility of calendar year reserve movements and the corresponding reserve sizes for these movements.
- **Bootstrap method** to assess volatility within historical claims triangles.

The bucket analysis (discussed in section 5) is our central approach, as it considers actual historical reserve movements and the volatility associated with them. The bootstrap approach (discussed in section 6) instead considers the volatility inherent within historical paid claims triangles without relying on subjective estimates of ultimate claims. The two methods have different strengths and weaknesses, which are outlined alongside the discussion of the analysis performed. Specific uncertainties and limitations of the bucket analysis are discussed in section 9.

We also analysed data from a reserving regulatory return from the Australian Prudential Regulation Authority (APRA). Only the bootstrap method was performed on this dataset, as the return includes insufficient datapoints to support the bucket analysis. The results of this analysis are discussed in section 6.6.

In addition, we analysed reserve risk data from LCP's annual capital benchmarking survey. The data from this are actual reserve risk CoVs selected by practitioners, and therefore reflect a different view than the historical volatility analyses performed on the other datasets. The results of this analysis are discussed in section 7.

Results from all analyses performed and our conclusions are outlined in section 8.

2.3. Scope

Reserve risk: Our focus has been on analysing reserve movements. The reason for this is that, in a capital modelling context, reserve risk is generally simpler to model than underwriting risk, as it is usually modelled as a single loss type, rather than being split between attritional, large and catastrophe losses. In particular, the skewed nature of natural catastrophe losses (which, due to the short tailed nature of their payment patterns, have less of an impact on reserve risk) may skew the results of any analysis on underwriting risk losses. Our view is that the results presented in this paper are likely also applicable to underwriting risk, but further work is needed to fully understand the extent to which this is the case.

Gross vs net: The analysis has been performed on a gross of reinsurance basis, to avoid confounding effects from different non-proportional reinsurance structures between firms.

Geography: The majority of the analysis has been performed on US data, since this is the source of the most extensive publicly available dataset. Results have also been sense checked where possible on a like-for-like basis against Australian data, although this dataset is much sparser and results are therefore less credible. Given that the focus of this research is on the fundamental nature of diversification within insurance portfolios, we have no reason to believe that the results would be different across territories.

Currency: Each of the three datasets used is based on a single currency: Schedule P is entirely in USD, APRA is entirely in AUD, and the LCP benchmarking surveys were converted to GBP based on the prevailing rates in force at the time the surveys were conducted. Since the benchmarking survey only includes data at a single point in time, we do not expect any impact on results from exchange rates changing over time. The objective of this exercise is to produce an estimate for a single dimensionless parameter, which should be invariant to currency effects. However, note that changes in capital model inputs due to movements in exchange rates would not normally be expected to result in changes to insurance risk volatility parameters, and the results of this exercise should only be applied to volumes that are already on a consistent exchange rate basis.

Basis: The techniques used cover both one-year and ultimate risk. The bucket analysis considers one-year historical reserve movements, the bootstrap considered is an ultimate bootstrap, and the survey data collected is on an ultimate basis. Since the focus of this research is on the fundamental nature of diversification within insurance portfolios, we have no reason to believe that different results would be obtained between ultimate and one-year views.

Our scope is limited to adjusting volatility for effects that are purely due to changes in volume, in line with the well-established principle that greater pooling of risk reduces volatility of outcomes, and vice versa (commonly known as the Law of Large Numbers). Accounting for changes in the nature of insurance business is more complex, and cannot necessarily be fully reflected by the approaches we outline here.

3. Approach

3.1. Model

We investigate modelling the relationship between volume (v) and volatility (CoV) of insurance losses (specifically, reserve movements) as a power curve described by the following equation:

$$CoV = av^{-b}$$

Equation 1

This relationship is shown in Figure 1 below for different values of b.

Figure 1

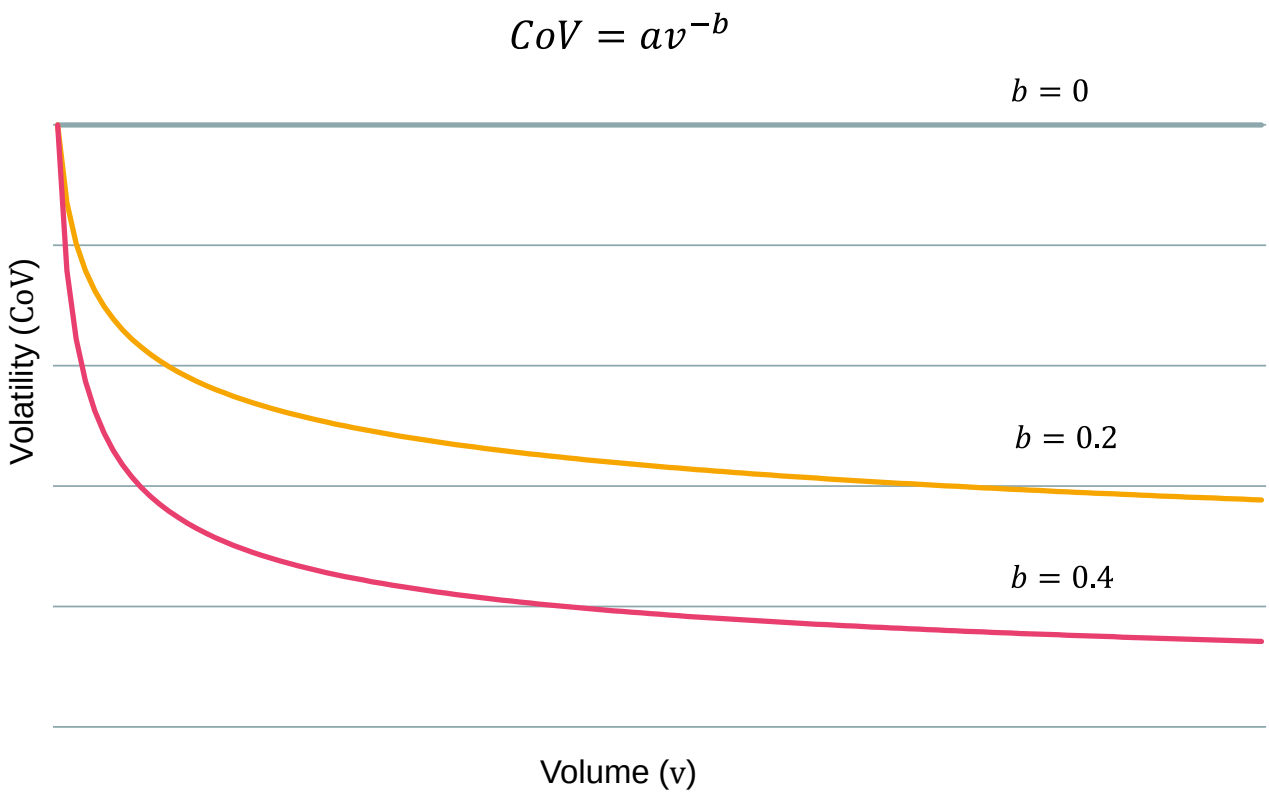

As shown in Figure 1, for $b > 0$, this model is consistent with the established principle that volatility decreases as volume increases, due to increased diversification. Higher values of b result in volatility being more sensitive to changes in volume. This indicates that the business has more specific risk, ie increasing volume results in greater diversification. Conversely, lower values of b imply that volatility is less sensitive to changes in volume. This is indicative of more systemic risk within the business that cannot be diversified away by increasing volume.

From this mathematical relationship linking volatility to volume, we can derive a volume adjustment factor that can be applied to a CoV to get a corresponding volume adjusted CoV. Define CoV_{before} and CoV_{after} as the original and volume adjusted CoVs respectively. Then from the above equation, we have

$$CoV_{before} = av_{before}^{-b}$$

Equation 2

$$CoV_{after} = av_{after}^{-b}$$

Equation 3

where v_{before}, v_{after} are, respectively, the original and subsequent non-zero volumes from which the CoVs have been calculated or estimated for. For example, CoV_{before} and v_{before} could be the CoV and volume based on experience at the previous year end, and CoV_{after} and v_{after} the parameterised CoV and volume projected for a future year-end.

Dividing Equation 3 by Equation 2 gives

$$\frac{CoV_{after}}{CoV_{before}} = \frac{av_{after}^{-b}}{av_{before}^{-b}}$$

$$\Rightarrow CoV_{after} = CoV_{before} * \left(\frac{v_{after}}{v_{before}}\right)^{-b}$$

Or equivalently,

$$CoV_{after} = CoV_{before} * \left(\frac{v_{before}}{v_{after}}\right)^{b}$$

Equation 4

The formula in Equation 4 implies that we can volume-adjust CoVs with a simple multiplicative factor $\left(\frac{v_{before}}{v_{after}}\right)^{b}$. The only information required to make the adjustment is the ratio of the volumes before and after, plus the parameter b.

Note that based on the above formula, provided $b > 0$, we have that:

If $v_{after} > v_{before}$ then $CoV_{after} < CoV_{before}$, and vice versa.

Intuitively, this makes sense, as if the volume of business has increased, then we would expect more diversification and therefore a lower volatility parameter.

Since the a parameter in the power curve model does not affect these calculations, the aim of this work is to parameterise b so that $CoV_{after} = CoV_{before} * \left(\frac{v_{before}}{v_{after}}\right)^{b}$

Equation 4 is a simple formula that can be used to volume-adjust CoVs, using only the ratio of volumes before and after.

We considered alternative possible relationships between volume and volatility to the power law model described above. In particular, we considered a model derived from probabilistic theoretical first principles, but found that the power law model provided a better fit to the data. Details of the theoretical model can be found in Appendix 4: Performance of theoretical model on real world data.

3.2. Bounds

We can derive upper and lower bounds on the possible values that b can take.

Lower bound: $b > 0$

A natural lower bound for b is 0, ie b cannot take negative values. A negative b would imply volatility increases as volume increases, which is not consistent with the pooling of risk principle.

The result $b = 0$ implies that volatility is completely invariant to changes in volume. That is, increasing the volume would result in no additional diversification benefit. This can be seen in the straight line in , above.

A relationship between volume and volatility of this nature (where $b = 0$) can be interpreted in the real world as an insurance portfolio entirely made up of systemic risk and no diversifiable risk. In practice, this situation is unlikely and therefore we are comfortable with a strict lower bound $b > 0$.

Upper bound: $b < 0.5$

For the upper bound, we show that, for a portfolio of N independent and identically distributed claims, we obtain a value of $b = 0.5$. Given that this scenario corresponds to a portfolio entirely made up of diversifiable risk and no systemic risk (as claim amounts are independent with each other and with N), this represents the maximum possible value of b, assuming that claim amounts cannot be negatively correlated with each other. Again, this situation is unlikely to occur in practice, resulting in a strict upper bound of $b < 0.5$.

Let X_i be the size of the i-th of N claims, where N is a random variable with distribution $N \sim Poisson(n)$

Assume the X_i are independent and identically distributed with mean μ and variance σ^2.

Then define S as the total claims:

$$S = \sum_{i=1}^{N} X_i$$

Then we have:

$$E[S] = E\left[\sum_{i=1}^{N} X_i\right]$$

Using the tower law of expectations, we can write this as

$$E\left[E[\textstyle\sum_{i=1}^{N} X_i \,|N]\right]$$

$$= E\left[N * E[X_i]\right]$$

$$= E[N] * E[X_i]$$

$$= n\mu$$

Similarly, using conditional variance theory, we have

$$Var(S) = E[Var(S|N)] + Var(E[S|N])$$

$$= E\left[Var\left(\sum_{i=1}^{N} X_i \,|N\right)\right] + Var(E[\textstyle\sum_{i=1}^{N} X_i \,|N])$$

Because the $X_i s$ are independent, this gives us

$$E[N * Var(X_i)] + Var(N * E[X_i])$$

$$= n\sigma^2 + E[X_i]^2 * Var(N)$$

$$= n\sigma^2 + \mu^2 n$$

Since we are looking at CoV as our volatility measure, we have:

$$CoV(S) = \frac{\sqrt{Var(S)}}{E(S)} = \frac{\sqrt{n}\sqrt{\sigma^2 + \mu^2}}{n\mu} = \frac{\sqrt{\sigma^2 + \mu^2}}{\mu\sqrt{n}}$$

Then defining volume v to be the expected value of the total claims S_n:

$$v = E[S] = n\mu$$

This results in:

$$CoV(S) = \frac{\sqrt{\sigma^2 + \mu^2}}{\mu\sqrt{n}} = \frac{\sqrt{\sigma^2 + \mu^2}}{\sqrt{\mu}} * \frac{1}{\sqrt{n\mu}} = av^{-0.5}$$

Equation 5

where constant $a = \frac{\sqrt{\sigma^2+\mu^2}}{\sqrt{\mu}}$.

This result can also be shown empirically to hold for a range of commonly used distributions for N and X_i, where the X_i are independent and identically distributed (and also independent of N).

We have shown that the parameter b is expected to take values between 0 and 0.5. A value of 0 implies a portfolio with only systemic risk, while 0.5 implies a portfolio with only specific risk. In practice, most insurance portfolios will contain a mixture of both specific and systemic risk, and therefore we expect an estimate for b that falls towards the middle of this range.

4. Data

This section provides a brief overview of the three datasets (Schedule P, APRA and LCP capital benchmarking survey) relied upon for the analyses performed. Some further details, in particular a description of any data cleaning that we carried out, are provided alongside the descriptions of the relevant analyses. Descriptions of class, duration etc are provided as we have attempted to fit models that consider these as factors, and also to demonstrate that the data analysed include a broad range that is likely representative of the market as a whole.

The APRA data has been provided to us under a Creative Commons Attribution 3.0 Australia Licence (CCBY 3.0), and we note that neither APRA nor the NAIC has endorsed the results presented in this paper.

4.1. Schedule P

The bulk of our analysis was performed on the Schedule P dataset from the National Association of Insurance Commissioners (NAIC). This data includes claims information for major personal and commercial lines for all property and casualty insurers that write business in the US. This dataset was used because it is large enough to provide credible results, and easy to analyse as the data are provided in a consistent format. The data used in this analysis include the 10 years over the period 2013-2022 and contains reserving metrics including Paid, Incurred and Ultimate claims for each entry of each reserving triangle.

The dataset has total reserves of around $1tn as at December 2022, spread across 13 reserving classes and 9,403 paid claims triangles. The chart below shows the reserves and number of triangles in each class. We have also grouped the classes into short, medium and long-tailed based on the claims development patterns. Short-tailed classes are those that reach substantially 100% developed in less than five years, medium-tailed classes are those which develop in between 5 and 8 years, and long-tailed classes take longer than 8 years to develop to ultimate.

Figure 2

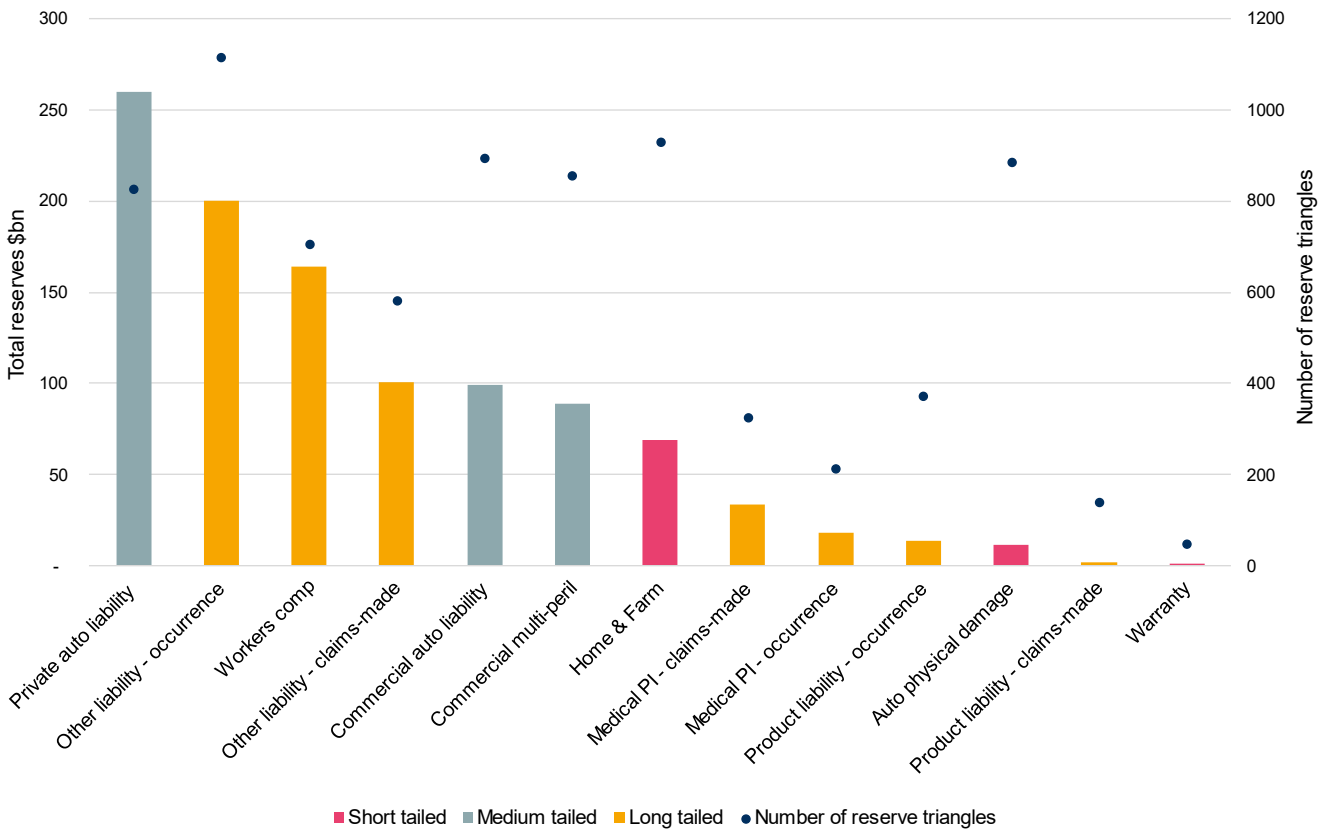

We can see that the dataset is skewed towards medium and long-tailed classes, with a much smaller proportion of total reserves in the short-tailed classes. However, the chart shows that the total number of reserving triangles is generally similar across short, medium and long-tailed classes.

4.2. APRA

Our secondary dataset to be analysed was a reserving regulatory return from the Australian Prudential Regulation Authority (APRA). This dataset is much smaller than the Schedule P dataset, and total reserves are less than AUD $50bn. There are 16 classes of business within this data. Figure 3 shows the breakdown of reserves by class, split into APRA's defined groupings of Short tailed, Long tailed and Other business.

Figure 3

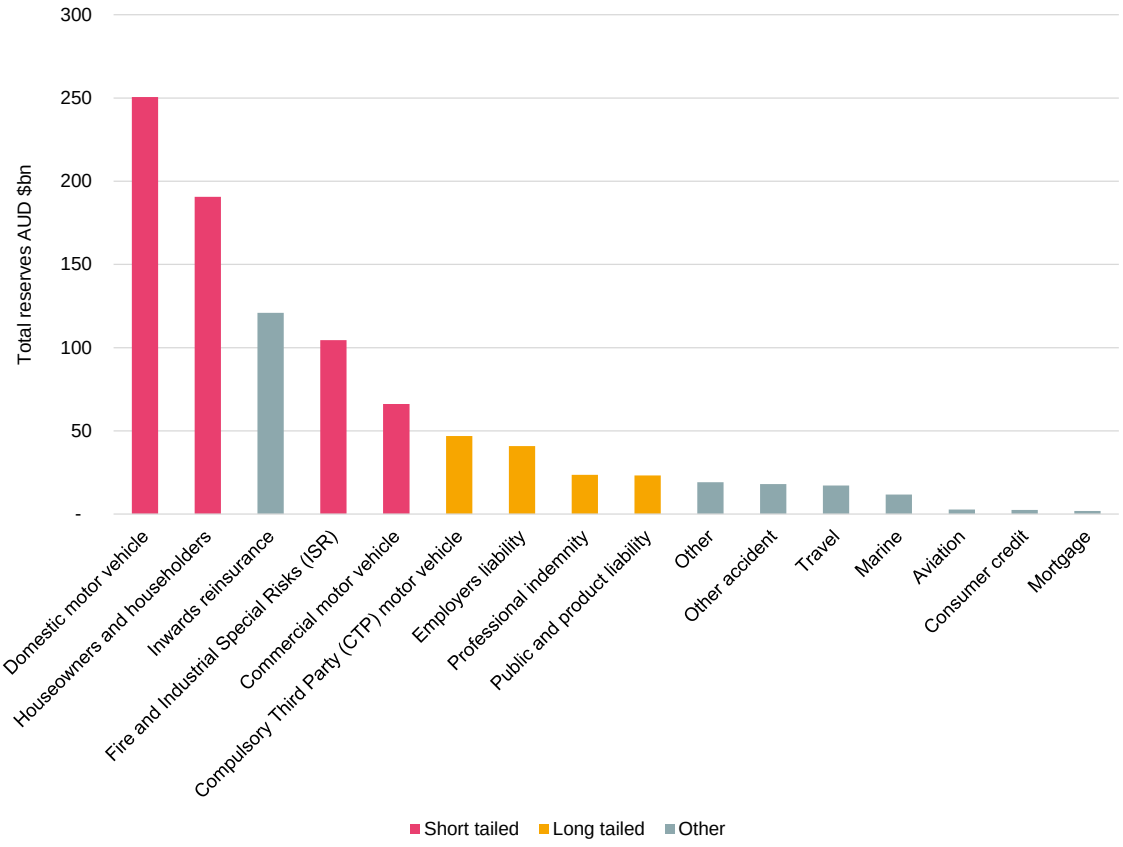

4.3. LCP capital benchmarking survey

We analysed the results from the 2023 and 2024 LCP capital benchmarking surveys. As part of the 2024 survey, we collected capital modelling data from 37 different companies across the London Market. The data from this survey was collected during April 2024 and a survey report was published in July 2024. The 2023 survey included data from 30 different companies across the London Market, with data collected during the first seven months of 2023 and was published in August 2023.

One aspect of the surveys involved analysing data on reserve volumes and CoVs for each survey respondent's classes of business. In order to compare classes of business, survey respondents were asked to assign each of their classes of business to one of the LCP survey classes, as set out below:

Property	Casualty	Other
Motor	A&H	Aviation
Property – Binder	Casualty – EL & PL	Cyber
Property – D&F	Casualty – FI & PI	Energy
Property – General	Casualty – General	Marine
Inwards reinsurance – Property	Inwards reinsurance – Casualty	Inwards reinsurance – Specialty
		All other classes

Each LCP survey class was categorised as one of Property, Casualty or Other as shown in the table.

Figure 4

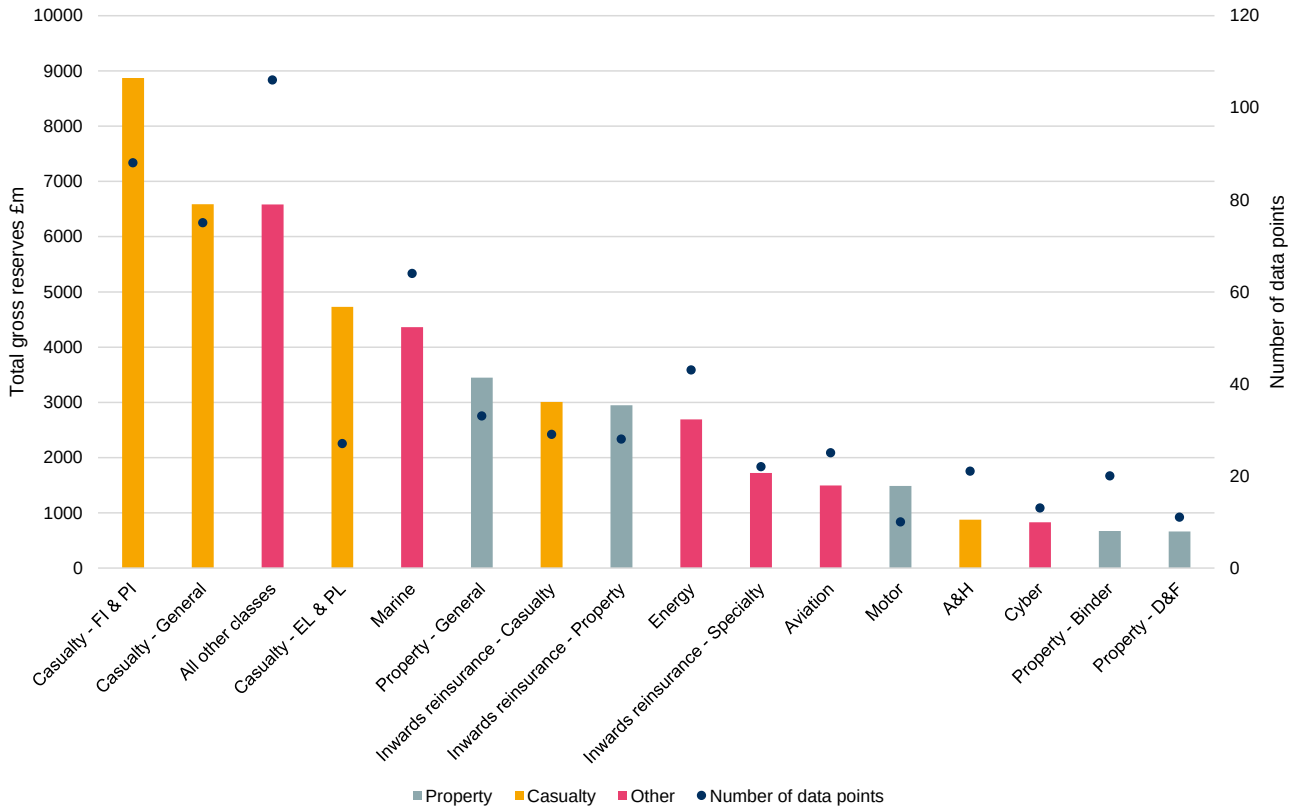

5. Bucket analysis

The bucket method analyses the relationship between volatility of calendar year reserve movements and the corresponding reserve sizes for these movements. The approach of bucketing data is important in this analysis, as to calculate reserve volatility, we need to apply a volatility measure to a set of data points. Therefore, to analyse how volatility changes with different reserve sizes, we need to create subsets of sufficiently similar data points.

5.1. Methodology

The bucket analysis involved performing the following steps on the Schedule P data (which is described in Section 4.1):

1. Calculate calendar year reserve movements

For each company and class of business, we calculated the % movement in reserves over a given calendar year $(t, t+1)$. We do not want to include the run-off of reserves or new years of account in the calculation, as this would overestimate the reserve volatility measure. To prevent this, we adjust the time $t+1$ reserves by adding back the paid claims over the year and removing the reserves associated with the new year of account.

$$Adjusted\ reserves_{t+1} = Reserves_{t+1} + (Paid_{t+1} - Paid_t)$$

where $Reserves_{t+1}$ denotes the reserves at the end of the calendar year t, excluding the most recent underwriting year.

We can summarise the whole calculation using the following formulae:

$$\%\ Reserve\ movement_t \coloneqq r_t = \frac{Adjusted\ reserves_{t+1}}{Reserves_t} - 1$$

$$= \frac{Reserves_{t+1} + (Paid_{t+1} - Paid_t)}{Reserves_t} - 1$$

$$= \frac{Ultimate_{t+1} - Paid_{t+1} + (Paid_{t+1} - Paid_t)}{Ultimate_t - Paid_t} - 1$$

$$= \frac{Ultimate_{t+1} - Paid_t}{Ultimate_t - Paid_t} - \frac{Ultimate_t - Paid_t}{Ultimate_t - Paid_t}$$

$$= \frac{Ultimate_{t+1} - Ultimate_t}{Ultimate_t - Paid_t}$$

where:

- r_t denotes the percentage reserve movement over calendar year t;
- $Adjusted\ reserves_{t+1}$ denotes the adjusted reserves at the end of calendar year t;
- $Reserves_t$ denotes the total reserves at the start of the calendar year t;
- $Ultimate_t$ denotes the total ultimate claims at the start of calendar year t; and
- $Paid_t$ denotes the total paid claims at the start of calendar year t.

Note that each calendar year reserve movement r_t has an associated reserve volume $Reserves_t$, which we will now denote as R_t.

Plotting R_t against r_t gives the following chart:

Figure 5

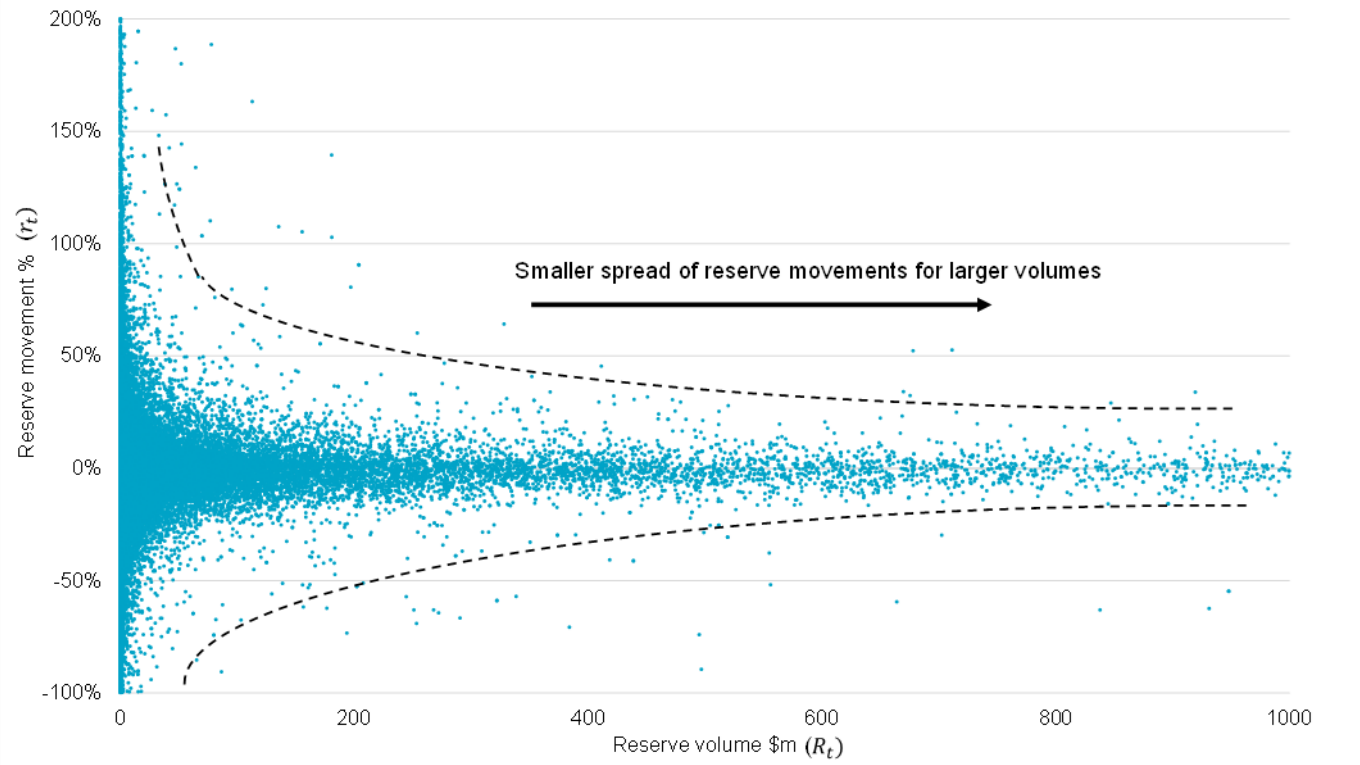

The x axis and y axis limits have been reduced to aid readability of the chart. As a result, note that there are reserve movements that do not fit onto the axis limits of the chart.

In Figure 5, each data point corresponds to one reserve movement r_t. In total, there are around 60,000 data points across all combinations of companies, classes of business and calendar years.

One notable feature of the chart is how the spread of reserve movements becomes smaller for larger reserve volumes. This is evident from the clear 'funnelling' shape of the graph.

2. Filter out extreme reserve movements

As part of cleaning the data, we filtered out reserve movements outside the range $r_t \in (-100\%, 500\%)$.

Note that $r_t < -100\%$ would imply that the reserves were positive at the start of the year and have become negative by the end of the year (or vice versa). Given that this type of reserve behaviour is unusual, we removed these movements from our analysis. Similarly, very large reserve movements $r_t > 500\%$ are likely either due to data artefacts, or material changes in risk profile, and therefore do not fit the profile of adjustments we are attempting to model for. In order to avoid these data points skewing the quantitative metrics used to measure volatility of reserve movements, we removed them from the dataset.

Figure 6

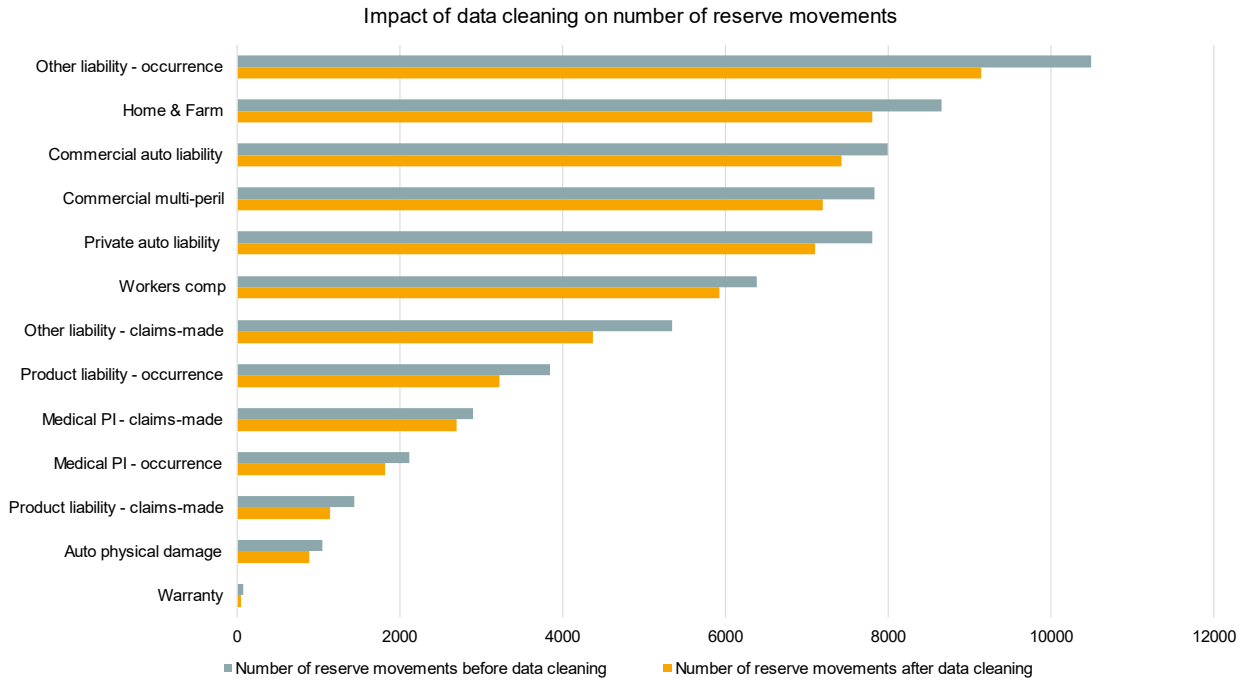

Figure 6 shows the impact of the data cleaning performed on the number of data points for each class of business within the Schedule P dataset. We can see that the proportional impact is broadly similar across all classes, ie no one class is disproportionately affected by data cleaning. Based on this and the analysis shown in Figure 9, we are satisfied that the data cleaning results in a dataset more comparable with the types of reserve movements we are attempting to model in practice.

3. **Data bucketing**

Now consider pairs of reserve movements (R_t, r_t), as shown in Figure 5.

We sorted these pairs in order of increasing reserve size such that $R_1 < R_2 < \cdots$, etc, to create a sequence

$$(R_1, r_1), \ldots, (R_i, r_i), \ldots,$$

Define the set of all re-indexed pairs of reserve movements as

$$\Gamma = \{(R_i, r_i)\}_i$$

We can then group the data into buckets $B_1, B_2, B_3, \ldots, B_j, \ldots$ using a bucket size n_B such that for all j:

$$(R_i, r_i) \in B_j \;\; \leftrightarrow \;\; (j-1) * n_B \; < i \leq j * n_B$$

That is, the first n_B reserve movements $(R_1, r_1), (R_2, r_2), \ldots, (R_{n_B}, r_{n_B})$ go into B_1, the next n_B reserve movements go into B_2, etc.

For this analysis we selected $n_B = 200$. Further details of this selection are shown in section 9.6.

Within each bucket B_j, we calculated the mean reserve size $\bar{R}_j$ and the volatility σ_j of the reserve movements within the bucket using the following formulae:

$$\bar{R}_j = \frac{1}{n_B} \sum_{i \in B_j} R_i$$

Equation 6

$$\sigma_j^2 = \frac{1}{n_B - 1} \sum_{i \in B_j} \left(r_i - \bar{r}_j\right)^2$$

Equation 7

At this stage, we have now calculated the mean reserve size $\bar{R}_j$ and volatility σ_j of each bucket. This results in a bucketed dataset:

$$\Sigma = \left\{\left(\bar{R}_j, \sigma_j\right)\right\}_j$$

Equation 8

We now want to understand the relationship between $\bar{R}_j$ and σ_j across the buckets. Referring back to Equation 1:

$$CoV = av^{-b}$$

We can map our variables in the following way to reconcile to this equation:

$$CoV = \sigma_j \qquad \text{and} \qquad v = \bar{R}_j$$

We can then plot these bucketed data points to get the following exhibit:

Figure 7

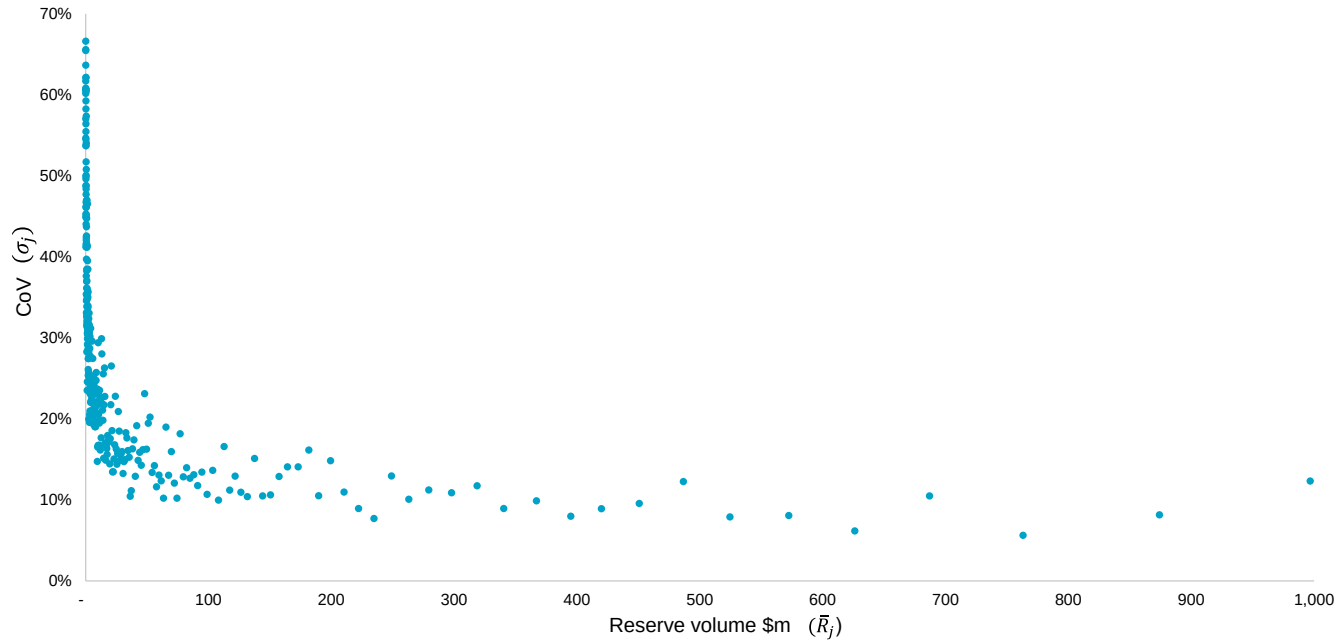

To analyse the relationship between volume and volatility further, we take the logarithms of Equation 1, as follows:

$$CoV = av^{-b}$$

$$\Rightarrow \log CoV = \log a - b * \log v$$

Plotting $\log CoV$ against $\log v$ gives the following graph:

Figure 8

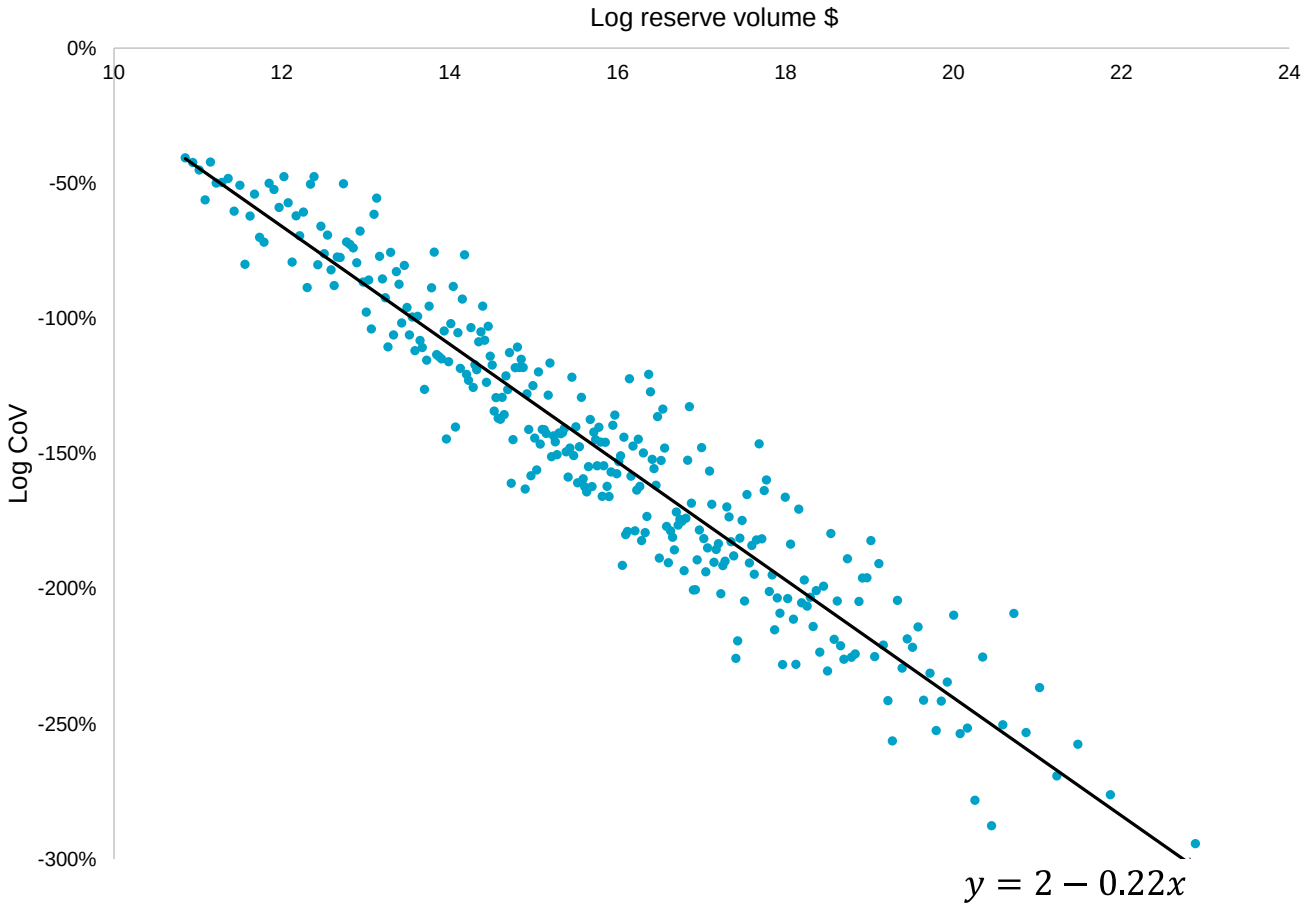

From Figure 8, we can see a linear relationship between $\log v$ and $\log CoV$. We used linear regression to fit a model to the data points shown above.

The fitted line of the model is shown in chart, and re-working to our original formula:

$$\log CoV = 2 - 0.22 * \log v$$

$$\Rightarrow CoV = av^{-0.22}$$

and therefore we obtain

$$\hat{b} = 0.22$$

To assess the goodness of fit for each model that we have fitted, we have used the R^2 metric, defined in Appendix 1: R-squared. The R^2 obtained for the model is 91%, implying that the proportion of unexplained variance in the model is very low.

Figure 9 below, shows how Figure 8 would look without the data cleaning step in the bucket analysis. Note that the two sets of points are separate, because the data have been re-bucketed before and after the exclusions were applied, which also impacts the estimates of volatility.

Figure 9

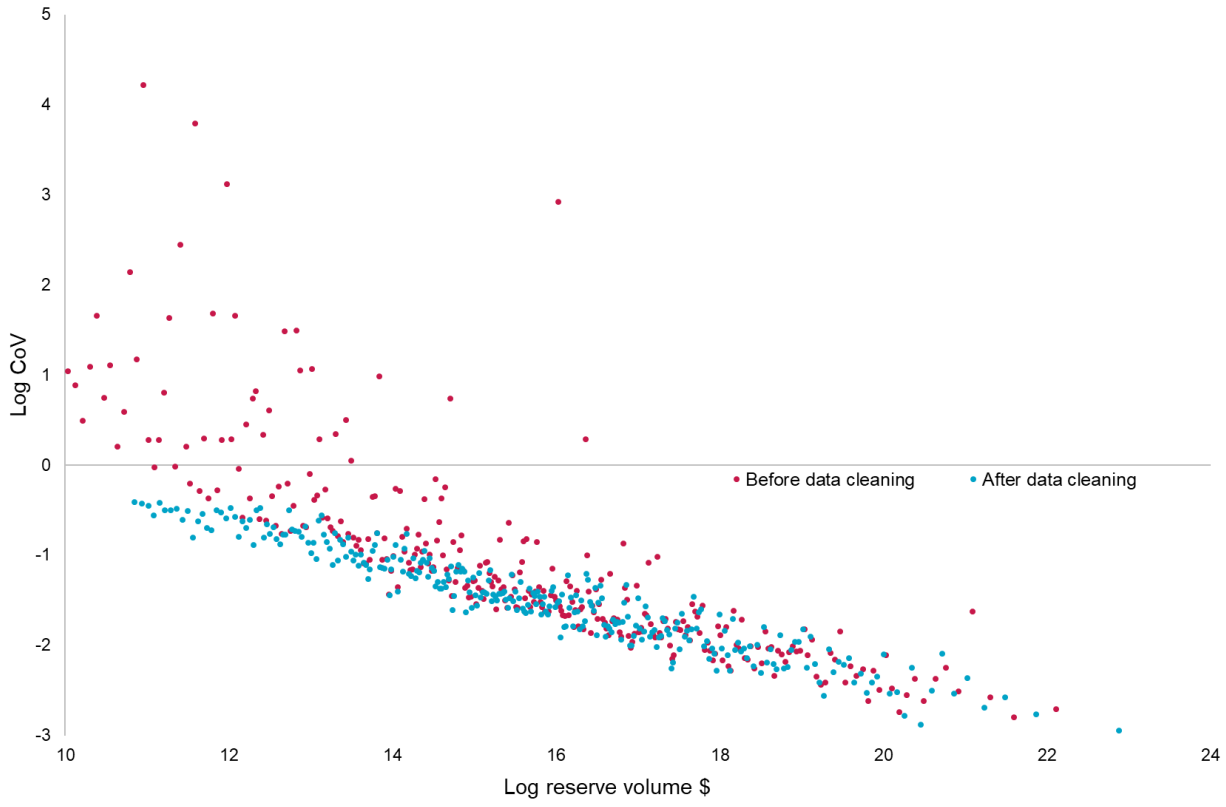

Given that the logarithm of a negative value is undefined, bucketed data points with a negative reserve volume (there are only two such points) have been excluded from the above chart. From the chart, we can see that for larger reserve volumes, there is only a small difference in vertical position of the bucketed data points before and after the data cleaning. However, for smaller reserve volumes, there is a much bigger difference. This shows that the data cleaning has a larger impact on buckets corresponding to smaller reserve volumes. For these buckets, very large reserve movements, which are likely data artefacts rather than genuine reserve deteriorations, skew the volatility upwards, leading to abnormally high CoVs. We therefore believe that it is appropriate to remove these large reserve movements given their significant influence on the vertical positioning of the bucketed data points.

5.2. Systemic volatility model

The above model $CoV = av^{-b}$ implies that

$$\lim_{v\to\infty} CoV = \lim_{v\to\infty} av^{-b} = 0$$

This means that as volume v grows large, the volatility of reserves tends towards 0.

Intuitively, we expect that as volume v grows large, the volatility of reserves tends towards some minimum value $\gamma > 0$, ie some proportion of the reserve volatility is undiversifiable. To allow for this feature, we considered another model:

$$CoV = av^{-b} + \gamma$$

Equation 9

With this new equation, we have

$$\lim_{v\to\infty} CoV = \lim_{v\to\infty} av^{-b} + \gamma = \gamma$$

And we also have

$$CoV > \gamma$$

for $v > 0$.

Equation 10

Finding the minimum volatility γ

We cannot fit a linear model to Equation 9 because the method of taking the logarithm of both sides results in:

$$\log CoV = \log(av^{-b} + \gamma)$$

which no longer simplifies into a linear relationship between $\log CoV$ and $\log v$. We therefore looked for an alternative approach to fitting this model. We can recover the linear relationship by subtracting the undiversifiable volatility which we call γ_0 and continue with our previous approach.

Given a value $\gamma = \gamma_0$, we define a new variable:

$$\widehat{CoV} = CoV - \gamma_0$$

Then we have:

$$\widehat{CoV} = av^{-b}$$

which has a similar structure to our original Equation 1, to which we can therefore fit a linear model, and subsequently calculate a model fit R^2.

Therefore, given a value $\gamma = \gamma_0$, we can fit a linear model $Model_0$ and corresponding model fit R_0^2. This allows us to map a value γ_0 to model fit R_0^2.

We can formalise this by defining a function f:

$$f(\gamma_0) = R_0^2$$

We then set up an optimisation problem to maximise the function f over different values of γ.

Doing so resulted in a maximum when $\hat{\gamma} = 0$, implying that the base model, which allows for no undiversifiable volatility γ, produces the best model fit.

As an alternative approach, we can set up an optimisation problem where we define:

$$g: x \to ax^{-b} + \gamma$$

And we try to minimise the objective function:

$$MSE = \sum_j \left(g(\bar{R}_j) - \sigma_j\right)^2$$

where $\bar{R}_j$ is the mean reserve size and σ_j is the volatility within bucket B_j.

We used a gradient descent algorithm to solve this optimisation problem, resulting in parameters

$$\hat{a} = 7.63 \qquad \hat{b} = 0.22 \qquad \hat{\gamma} = 0.00003$$

Note that the $\hat{a}$ and $\hat{b}$ parameters are the same as we obtained for the original model, and the fitted $\hat{\gamma}$ is very close to zero.

Commentary

Note that if Equation 9 held and there was undiversifiable volatility γ, then we would have $\log \gamma$ as a strict lower bound to the graph. In other words, we would expect the graph to start levelling off horizontally for large enough values of v. Looking at the chart below, we can support the results from the above analysis implying $\hat{\gamma} = 0$. We see that there is no evidence of the bucketed data points levelling off horizontally for large values of v, and therefore we have graphical evidence that there is no undiversifiable volatility in Schedule P. In addition, the chart shows the shape of the models we would expect to see if there was undiversifiable volatility for parameters $\gamma = 0.05, 0.1$ and 0.15 respectively. Note that for each undiversifiable volatility model shown below, we have refitted $\hat{a}$ and $\hat{b}$ to obtain the best model fit for that particular value of γ.

Figure 10

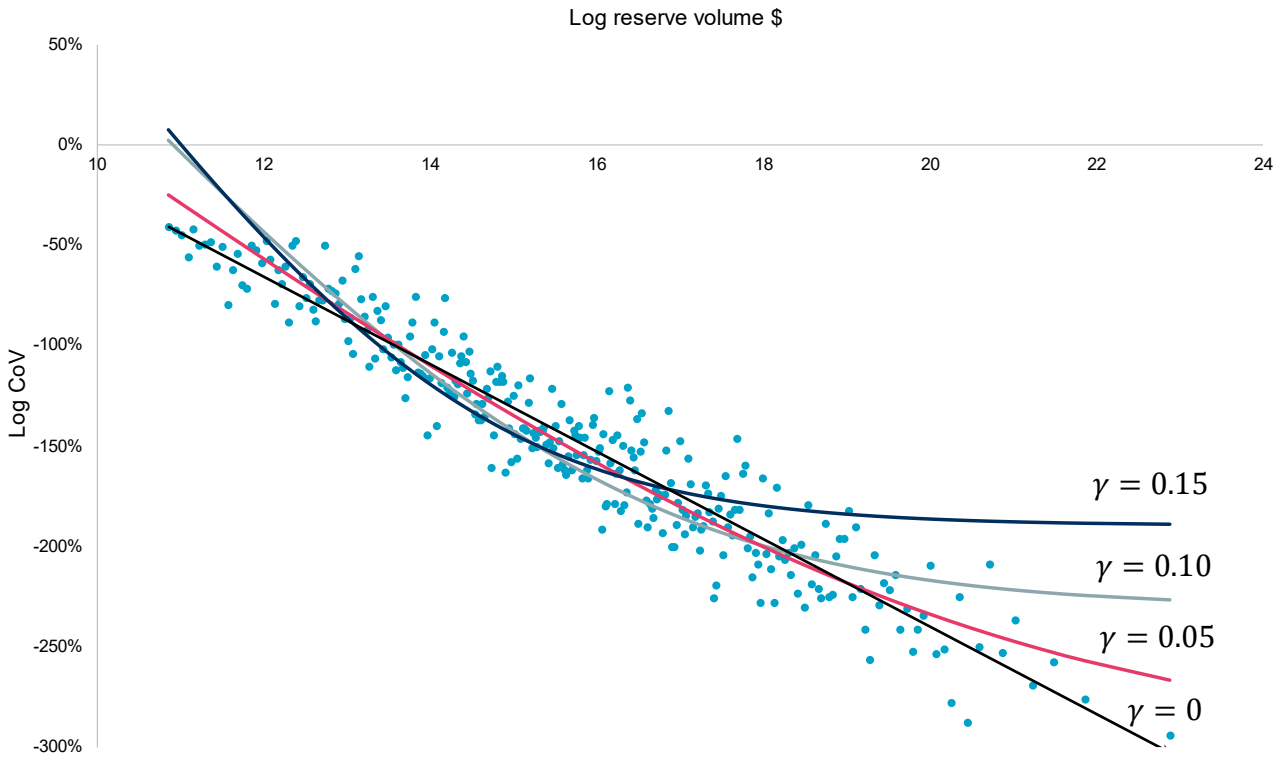

Note that the bucket with the largest reserve volumes (furthermost right data point on the graph) has a mean reserve size of $8.7bn, so one limitation of this analysis is that we are only analysing reserves up to this magnitude. It could be that this levelling off happens at larger reserve sizes, and we are just not able to see the levelling off by looking at this particular data set. However, in practice, there will be few insurance portfolios with homogenous class of business level reserves in excess of this level, so this limitation is not material.

5.3. Other factors: Class and duration

We also investigated the influence of class of business and tail length of the class on the relationship between volume and volatility. We expected different classes to have different balances of systemic (undiversifiable) and specific (diversifiable) risk and so different relationships with, or sensitivity of volume to, volatility.

We fitted models to the Schedule P data allowing for specific parameters for class of business and tail length. This was to answer the following questions:

- How class dependent is the relationship between volume and volatility?
- Is the volume volatility relationship driven by the length of tail of a given class of business?

- Does allowing for class or tail specific parameters materially improve the fit of the model?

Methodology

The following additional models were fitted:

Models	
Class model (regular)	$CoV = a_k v^{-b}$
Class model (advanced)	$CoV = a_k v^{-b_k}$
Duration model (regular)	$CoV = a_l v^{-b}$
Duration model (advanced)	$CoV = a_l v^{-b_l}$

for $k \in$ class of business, and $l \in$ duration group.

The Class model (regular) fits different a parameters for each class of business and fits a universal b parameter that is class independent. The Class model (advanced) fits different a and b parameters for each class of business. The classes are as specified in the Schedule P dataset.

Similarly, the Duration model (regular) fits different a parameters for each duration type and fits a universal b parameter that is independent of duration. The Duration model (advanced) fits different a and b parameters for each duration type. The classes are as specified in the Schedule P dataset. The duration types – Short, Medium and Long – were assigned to each class as described in section 4.1 and as shown in Figure 2.

Bucketing

This process is described for the class models. A similar process was applied using l instead of k, which allows us to obtain duration-specific parameters for each duration type l.

We partitioned the data into the respective classes of business. We then performed bucketing on each class of business separately in the same way as set out in step 3 in section 5.1. We used the same bucket size $n_B = 200$ as the original bucketing method. Following the re-bucketing, each bucket had the same number of datapoints (R_t, r_t) as previously, but now contained only one class of business.

However, there is a trade-off between homogeneity in class vs reserve size with our two methods.

Bucketing at a total level means that buckets are class heterogeneous, that is, each bucket can contain reserve movements in relation to multiple different classes. As mentioned above, each class of business contains a different balance of systemic and diversifiable risk, which then impacts the size of reserve movements for a given reserve size.

Bucketing at a class level means that buckets are now class homogeneous, but there is less homogeneity in the reserve sizes within each bucket. By performing bucketing on only one class of business but using the same bucket size n_B, there will be a wider range of reserve sizes within each bucket.

Figure 11, below, is a plot showing the bucket widths for each bucket for both the total model and the class model. We can see that, as evidenced above, the buckets are much wider for the class model. Note that the bucket widths for the class model appear to be more volatile. This happens because we are performing the bucketing separately for each class and then combining the buckets into an overall dataset, so the bucket widths displayed correspond to different classes.

Figure 11

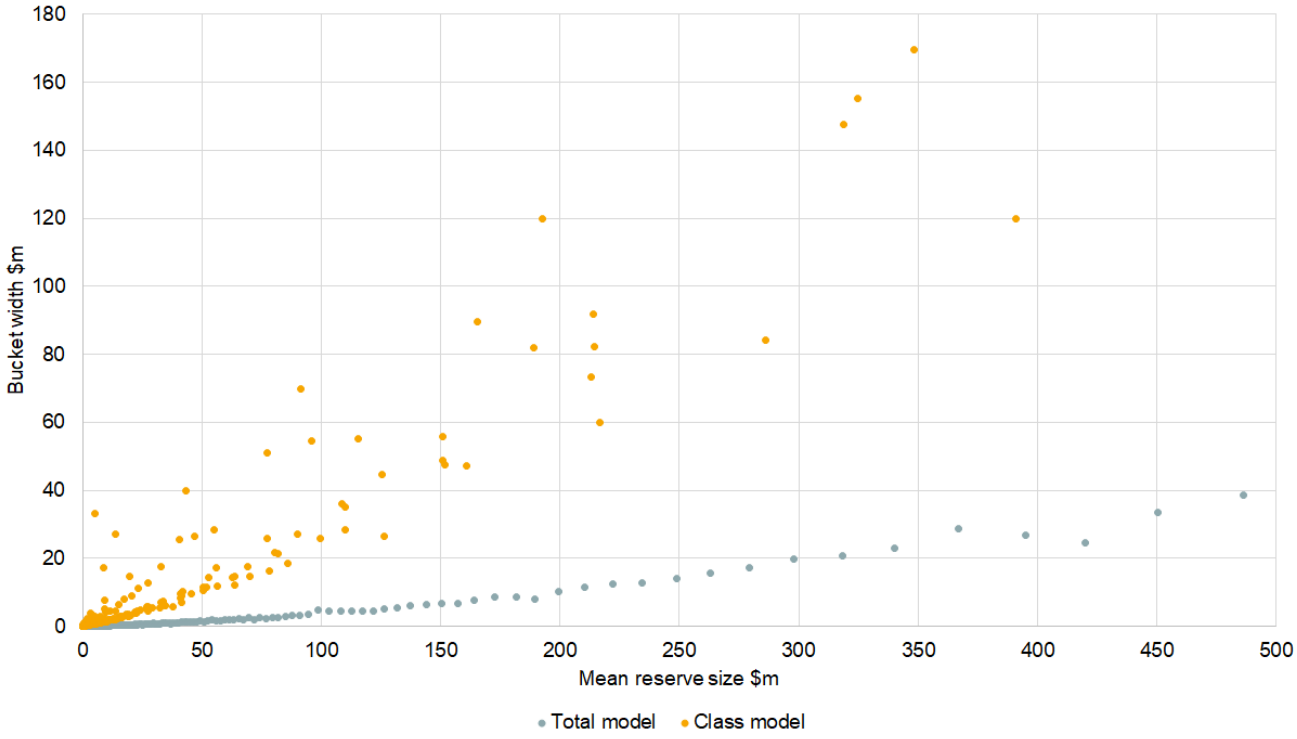

For each class k and bucket B_{kj}, we calculated the mean reserve size $\bar{R}_{kj}$ and volatility σ_{kj} of the reserve movements as before with Equation 6 and Equation 7. This results in a dataset of bucketed points for each class:

$$\Sigma_{\mathrm{k}} = \left\{\left(\bar{R}_{kj}, \sigma_{kj}\right)\right\}_j$$

We can then combine the class specific bucketed datasets to obtain an overall dataset of pairs of mean reserve size and volatility:

$$\Sigma' = \bigcup_k \Sigma_k$$

This new dataset has the same number of bucketed datapoints as the original bucketed dataset but with the addition of class labels, ie

$$|\Sigma'| = |\Sigma|$$

Or:

$$n_{\Sigma'} = n_{\Sigma}$$

where Σ is defined in Equation 8.

Mapping σ to CoV and R to v as before, we look at variations of our original Equation 1:

$$CoV = a_k v^{-b} \quad \text{and} \quad CoV = a_k v^{-b_k}$$

We again calculated the logarithms of CoV and v and fit linear models of the form:

$$\log CoV = \log a_k - b * \log v \quad \text{and} \quad \log CoV = \log a_k - b_k * \log v$$

This format allows us to obtain class-specific parameters for each class k.

Results

The table below are the results of the fitted parameters for the class and tail models, with the base model results repeated for reference:

Model		$\widehat{b}$	R^2	# parameters
Base model	$CoV = av^{-b}$	0.22	91%	2
Class model (regular)	$CoV = a_k v^{-b}$	0.22	90%	12
Class model (advanced)	$CoV = a_k v^{-b_k}$	0.12 - 0.28	60%	26
Duration model (regular)	$CoV = a_l v^{-b}$	0.22	91%	4
Duration model (advanced)	$CoV = a_l v^{-b_l}$	0.18 - 0.23	90%	6

Observations

Class model (regular): $CoV = a_k v^{-b}$

This model fits different a parameters for each class of business and fits a universal b parameter that is class independent. We obtained the same value $b = 0.22$ as our base analysis, with a similar R^2 value of 90%.

One interesting point of note is that the R^2 value for this model is slightly lower than the base model, despite introducing 10 new parameters. The bucketed datasets, Σ and Σ', underlying the two models are different, as the bucketing is on a total level and class of business level for the two models respectively. The lower R^2 value here is an example of the trade-off between homogeneity of class and reserve size as discussed earlier in this section.

Class model (advanced): $CoV = a_k v^{-b_k}$

This model fits different a and b parameters for each class of business. Note that we now have a range of b parameters across different classes of business, with 0.12 and 0.28 representing the two end points of the range. The R^2 value for this model is lower at 60%. This is mainly due to the higher level of parameter uncertainty around the estimation of the class specific b parameters.

Using class as a predictor of both a and b parameters is the equivalent of independently fitting a line of best fit through the bucketed data for each class of business separately, ie fitting a model to each dataset Σ_k. Bucketing the data points divides the total number of datapoints to fit a model through by bucket size $n_B = 200$, and therefore some classes do not have enough data points to produce credible estimates for the intercept and slope of the line of best fit through the bucketed points.

Duration models (regular and advanced): $CoV = a_l v^{-b}$ and $CoV = a_l v^{-b_l}$

These models work in a similar way to the class models, but for duration type rather than class. The regular model again produces the same value $b = 0.22$ as our base analysis, with a similar R^2 value of 91%. The advanced model produces a range of b parameters across different classes of business, with 0.18 and 0.23 representing the two end points of the range. As indicated by Figure 12, below, short tailed classes tend to produce lower values of b, but long tailed classes do not necessarily produce the highest values of b. Our conclusion is that class is more of a driver of b than duration, although there is a natural relationship between class and duration that results in an indirect relationship between duration and b.

Figure 12 shows the values of b produced by the Class model (advanced) for each class of business, with a vertical line showing the base model result of $\hat{b} = 0.22$.

Figure 12

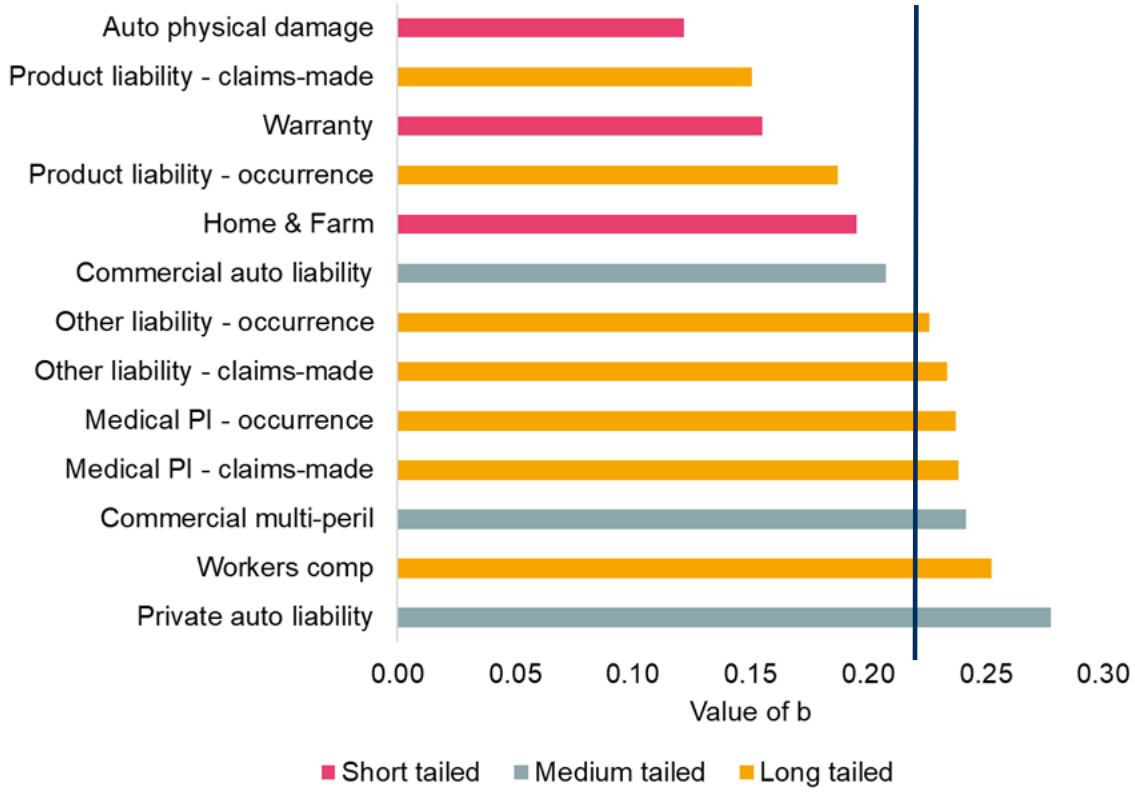

What causes the variation between b values across classes of business?

As discussed previously, the value of b in the relationship $CoV = av^{-b}$ can be thought of as a measure of the degree of sensitivity of volatility to changes in volume. In order words, it measures how much the business diversifies when volume increases. We can therefore think of class specific b values in terms of the level of systemic and diversifiable risk. Note that a larger value of b implies more sensitivity of volatility to changes in volume, and therefore a higher level of diversifiable risk.

The classes with notably lower values of b are the classes at the top of the chart in Figure 12. This implies that these classes are less diversifiable. For example, Auto physical damage relates to motor vehicle damage business, for which the claims profile is typically a large number of very small claims. The main causes of reserve movements for this class are likely to be systemic drivers such as inflation and less likely to be individual claims movements. Therefore, increasing the reserve volume of this class, (ie adding more small individual claims) is unlikely to greatly increase the level of diversification, ie volatility is less sensitive to changes in volume, which is consistent with the interpretation of a low b value.

Conversely, Workers' comp and Private auto liability have the largest values for b. This implies that these classes are more diversifiable. This may be because reserve movements for these classes are relatively more likely to be due to individual claims movements and relatively less likely to be systemic drivers. However, more work is needed to fully understand the impact of class (and duration) on the results we have obtained.

6. Bootstrapping analysis

The second part of our analysis employed a bootstrap method to assess volatility within claims triangles. This is a widely used method for calculating volatility in claims triangles.

It is a very different technique to measure volatility than the bucket analysis method, as it does not rely on subjective estimates of ultimate claims. As we can use the same dataset, we can compare the results of the bootstrap and bucket analyses and determine whether both methods yield consistent results despite the different techniques. This will help validate our findings and ensure conclusions are not biased by the method used.

6.1. Data

We filtered the Schedule P data on entities with greater than $1m reserves. This resulted in 5,666 triangles to bootstrap. This was done because some of the triangles in Schedule P are sparse, so filtering on larger companies provides better stability in the results.

6.2. Methodology

Bootstrapping is a resampling technique which we use to estimate the distribution of reserves from which a CoV can be derived. We used an Over-dispersed Poisson, or ODP bootstrap framework. See Appendix 2: Over-dispersed Poisson distribution framework for more detail on this distribution. This process can be summarised in 5 steps:

1. Fit a generalised linear model (GLM) with a log-link function and ODP error distribution to the incremental claims data. This is identical to fitting a chain ladder to the triangle.

2. Calculate the residuals from the fitted model.

3. Resample (with replacement) the residuals to create a pseudo triangle.

4. Refit the GLM using the pseudo-dataset to produce a new estimate of the reserves.

5. Repeat steps 3 and 4 n times to derive a forecast output for each pseudo-dataset. This gives a distribution of reserve estimates.

We selected $n = 10{,}000$. See Appendix 3: Bootstrapping methodology for more information on this process.

We therefore had, for every cohort in every reserving triangle, 10,000 reserve estimates. We calculated a CoV and corresponding mean reserve size for each.

We took logs of the CoVs and reserve estimates from the bootstrap output and fit a linear model. We call this the bootstrap base model. This model is of the structure of our original formula given in Equation 1:

$$CoV = av^{-b}$$

We also fit other models using features in the data as predictors. These include the Class model, Cohort model, and Class and cohort model. These are shown in the table below.

Models	
Bootstrap base model	$CoV = av^{-b}$
Class model	$CoV = a_k v^{-b}$
Cohort model	$CoV = a_l v^{-b}$
Class and cohort model	$CoV = a_{kl} v^{-b}$

for $k \in$ class of business and $l \in$ cohort

Each of these models have a universal b parameter, while the a parameters vary by class and/or cohort. The Class model is as before, with a different a_k for each class. The cohort model arises because we have a CoV and reserve estimate for each cohort in the triangle. We can therefore fit a model allowing for cohort to be a predictor, and so have a different a_l for each cohort in the triangle.

6.3. Benefits, uncertainty and limitations

The main benefit to this part of the work is that we are able to use a very different technique to measure the volatility, but using the same data as used for the bucket analysis. It is a purely data driven method that does not rely on ultimate claims estimates, which are by nature subjective and prone to biases.

However, as for all bootstrapping, there are a number of well-known limitations, including the fact that it is computationally intensive. We used LCP InsurSight's[2] reserve risk tool to handle the large quantities of data and calculations required.

Bootstrapping is not guaranteed to converge to the 'true' value, and we therefore filtered on entities with larger reserve sizes to provide greater stability in the results. Many of the other standard limitations of bootstrapping are partially mitigated by the fact that we are not using the estimated CoVs directly, but mapping them against volume to establish a relationship.

There are some additional limitations for this exercise. Because this is an automated approach, we did not apply any of the traditional actuarial judgements that would be applied as part of a standard bootstrapping exercise. This means that – for example – we did not smooth the data, exclude outlier link ratios, or apply any tail factors. Some of the triangles include unusual development patterns. This may affect the analysis, since one known limitation of the bootstrap is that a single unusual year of development (whether genuine or due to data artefacts) can result in a very high CoV.

[2] https://insursight.lcp.com/

6.4. Results

Figure 13

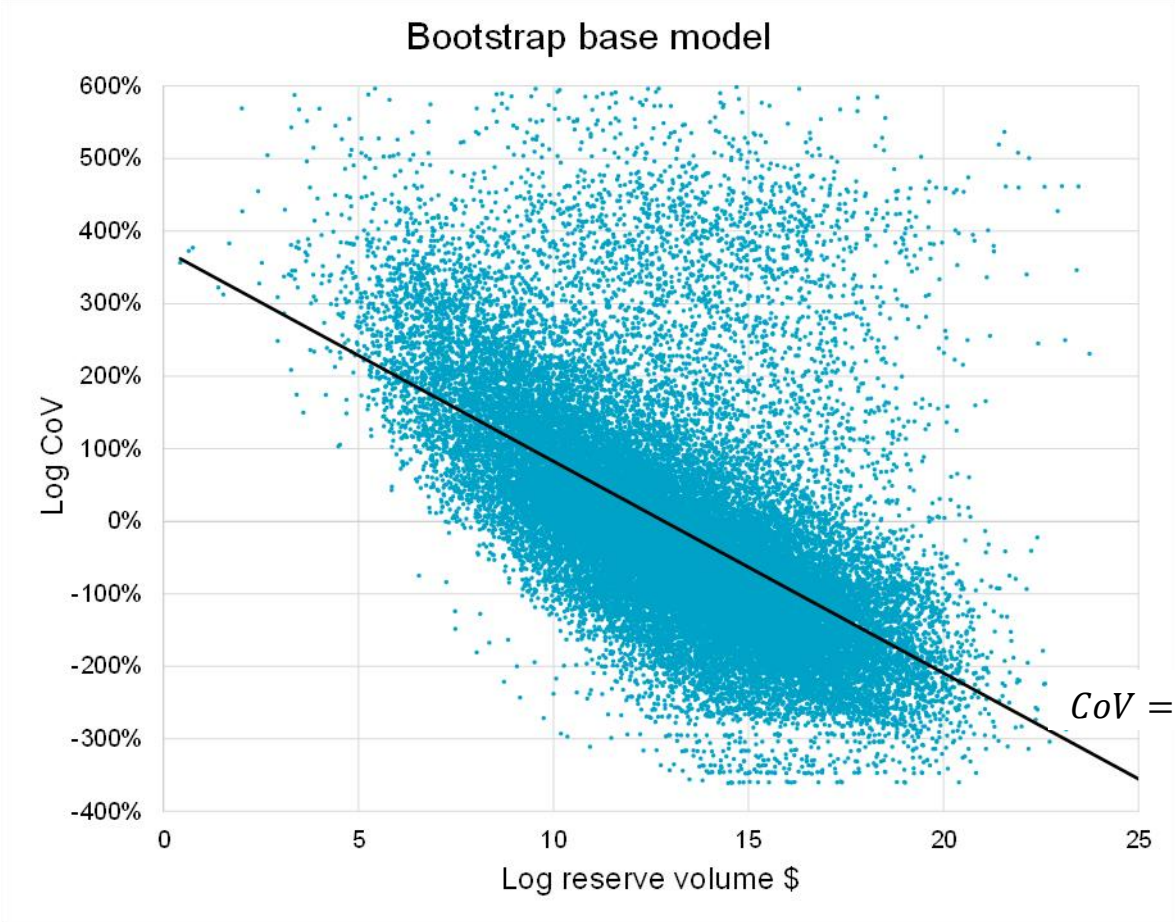

Figure 13 shows the bootstrap base model, ie the linear model fit between log CoVs and log reserve volumes. While there is a concentration of the data around the trendline, the sparser outliers can be more easily seen.

The bootstrap base model estimated a $\hat{b}$ parameter of 0.29 with an R-squared value of 33%. The relatively low R-squared is caused by the variance in the data not being explained by the model. This is due to the outliers seen in the chart. These outliers are likely caused by the inclusion of triangles with unusual development patterns, as discussed previously. Therefore, the relatively poor fit is not a material concern, in particular considering the low standard error of the parameter estimate.

The results of the other models, including the estimated values for $\hat{b}$, the associated R-squared value, and the calculated standard error associated with the estimate, are shown in the table below.

Model		$\hat{b}$	$SE(\hat{b})$	R^2	# parameters
Bootstrap base model	$CoV = av^{-b}$	0.29	0.002	33%	2
Class model	$CoV = a_i v^{-b}$	0.28	0.002	40%	12
Cohort model	$CoV = a_j v^{-b}$	0.26	0.003	35%	3
Class and cohort model	$CoV = a_{ij} v^{-b}$	0.25	0.002	42%	13

It is interesting to note that the fit generally improves by the introduction of parameters in the model (ie the R^2 value increases, albeit slightly), and the $\hat{b}$ estimate also gets slightly closer to the $b = 0.22$ obtained in the bucket analysis.

6.5. Bucketing the bootstrap

In order to mitigate the relatively poor fit observed in the bootstrap base model, we then combined the bucketing technique with the bootstrap approach. We grouped the bootstrapping output into buckets by reserve size. We again used buckets of size $n_B = 200$, consistent with the earlier analysis. We then took the median log reserve and median log CoV for each bucket. We chose to use the median because the mean was skewed by the extreme outliers we have seen, as discussed in section 5.1.

Fitting this model obtained $\hat{b} = 0.21$, which is very close to the base model estimate of 0.22. Bucketing also improved the fit over the pure bootstrap model, achieving an R-squared value of 95%. This is due to bucketing reducing the noise in the dataset.

Figure 14 below shows the fitted line of the bucketed bootstrap data:

Figure 14

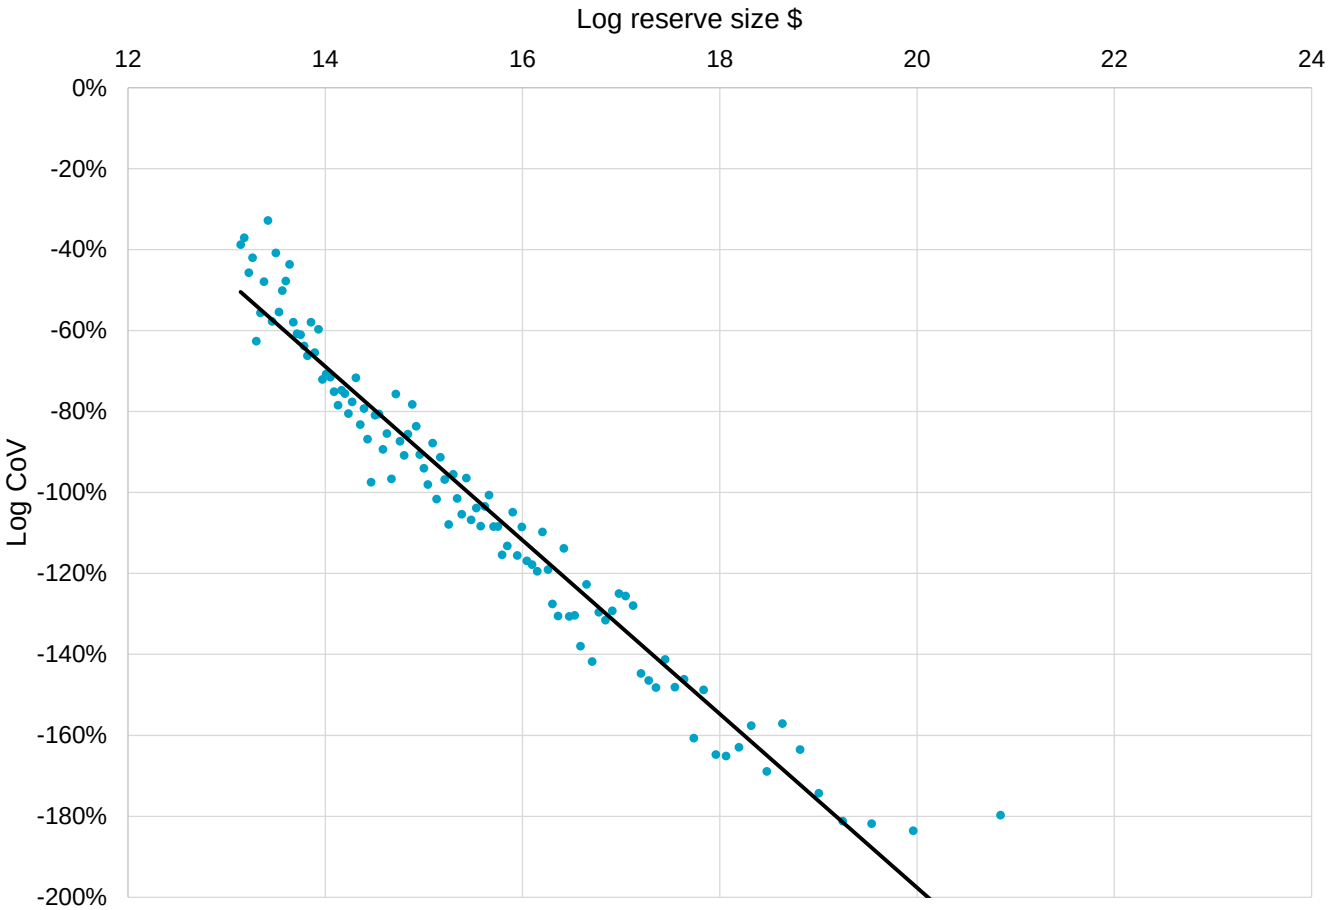

The datapoints are much tighter around the trend line, visually evidencing the improved fit. Note that in the chart, we can see some indication of a 'levelling-off' at higher reserve sizes. This could be evidence of the systemic risk that we discussed earlier in section 5.2.

6.6. APRA

We also looked at APRA data, which is described in Section 4.2. The data is sparser than Schedule P; however, bootstrapping the triangles and fitting a model obtains R-squared of 66%, which is a better fit than for Schedule P. The $\hat{b}$ parameter fitted is 0.31, higher than Schedule P models, suggesting there is a slightly greater sensitivity to volume on volatility for this data.

The bootstrap base model for Schedule P resulted in a $\hat{b}$ parameter of 0.29, which is broadly comparable to the result obtained for APRA. As these two datasets represent different regions, this provides comfort that the results are replicable and that the reliance on US data for the main analysis is not a material limitation. However, we note that the APRA dataset is fairly sparse, and therefore the results do not have high credibility.

The bucket analysis was not performed for the APRA dataset, as the data is too sparse to produce buckets of meaningful size, and the resulting estimates of volatility would therefore not be credible.

Figure 15

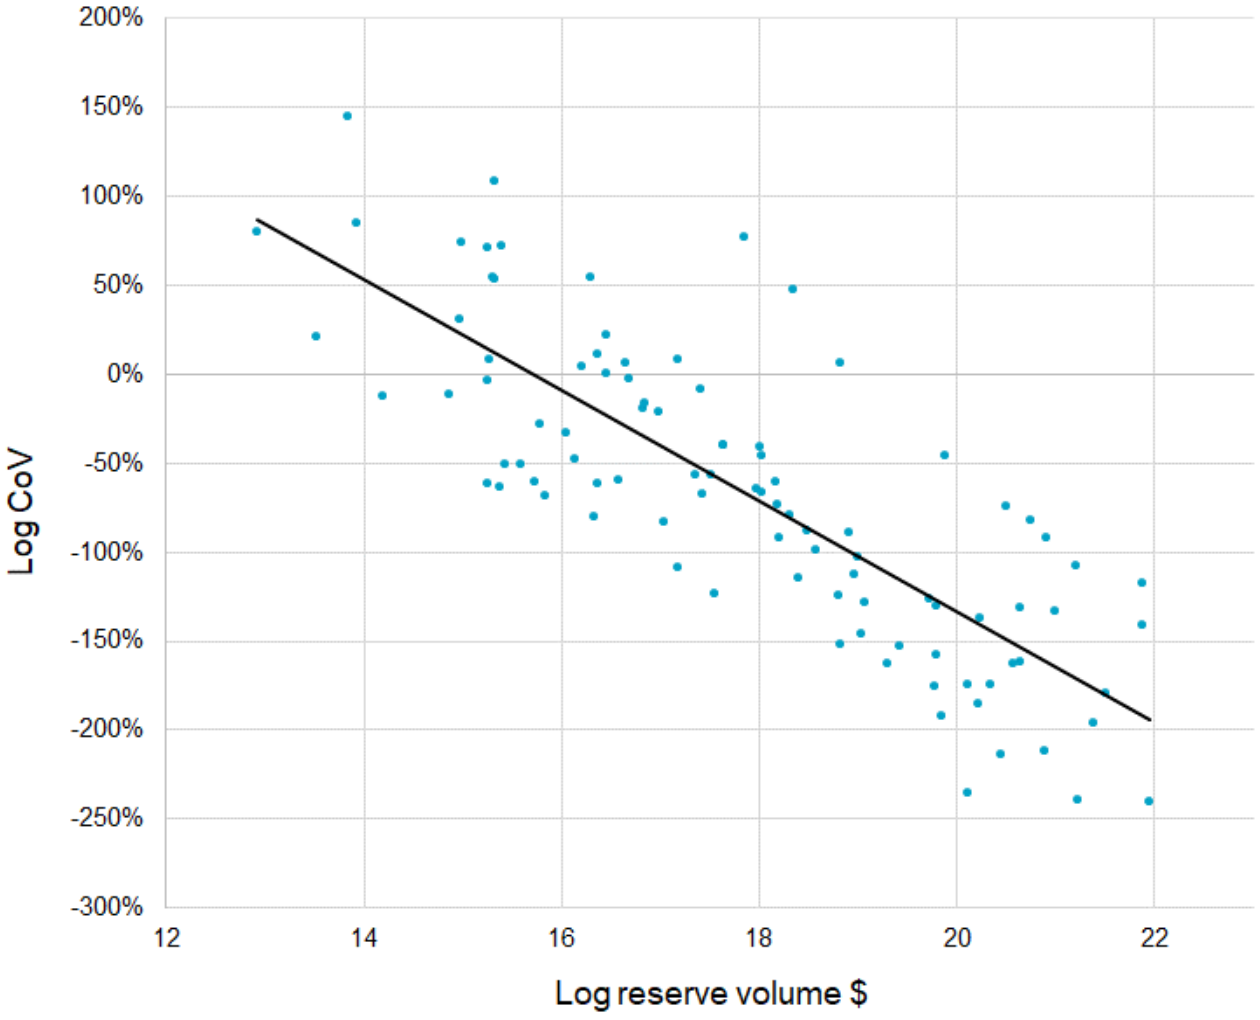

7. LCP capital benchmarking survey

Results from LCP capital benchmarking surveys (described in Section 4.3) provides an additional way of analysing the relationship between volume and volatility. The following data items were collected from participants for the survey analysis:

- Gross reserve volume by class of business
- Gross ultimate reserve risk CoV for the class as a whole

We can map these data items to volume and volatility respectively and perform a similar analysis as before. No bucketing or bootstrapping is required, as the required fields of volume and volatility are directly available without performing any additional calculations.

The following example slide is taken from the 2024 LCP capital benchmarking survey report:

Figure 16

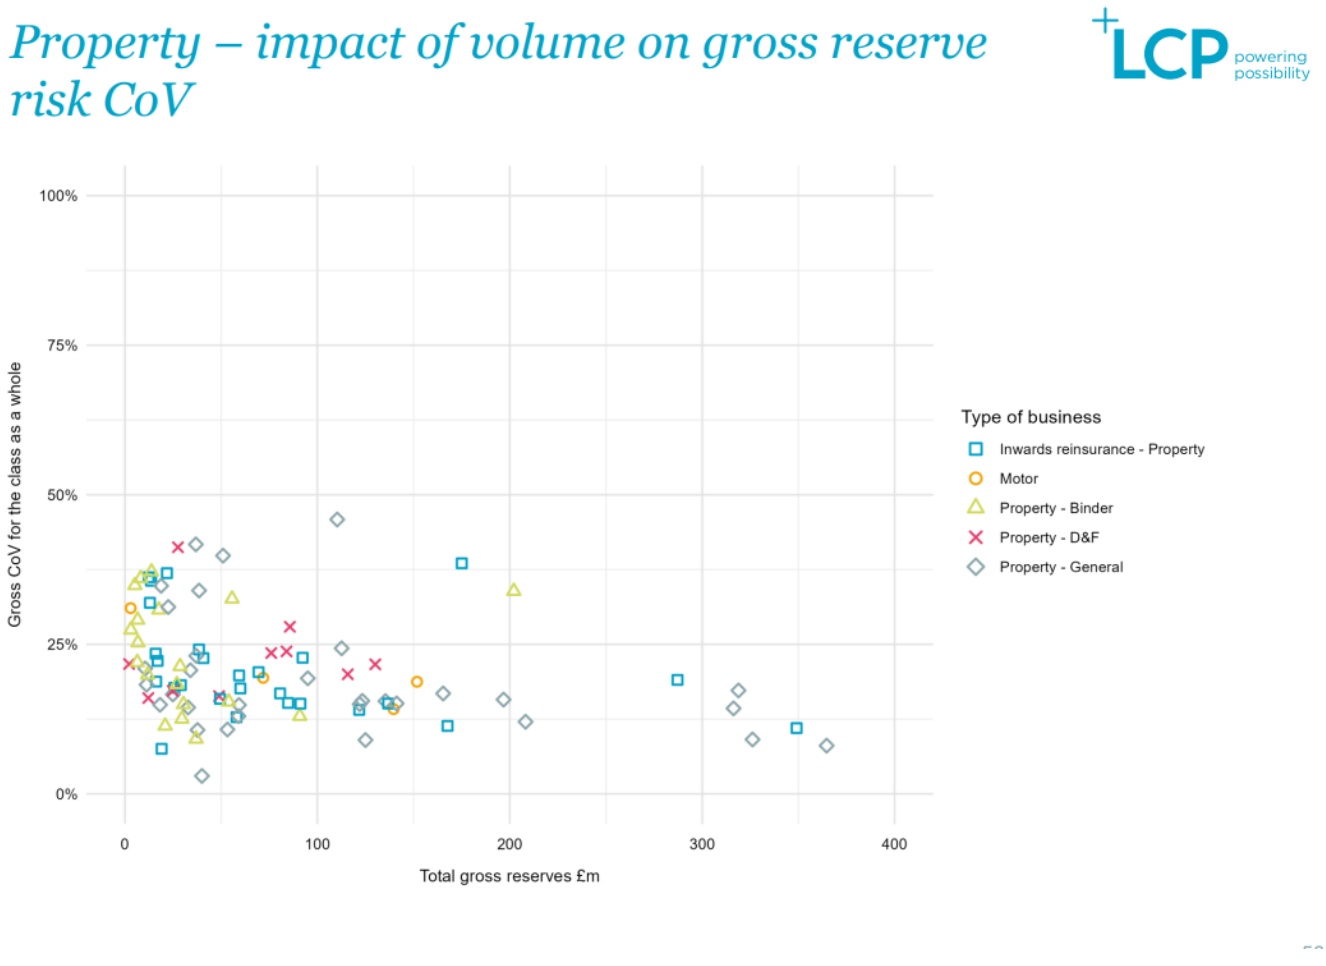

Figure 16 is a plot of the gross CoV for the class as a whole against the total gross reserves for each survey respondent's classes of business within the 'Property' category. Note that there is a no obvious power law relationship that we observed with the bucketing and bootstrapping analyses for Schedule P and APRA data.

We fit a linear model of the form $CoV = av^{-b}$ to the LCP capital benchmarking survey data for both 2023 and 2024 survey results.

Fitting a model to the 2024 benchmarking data resulted in a $\hat{b}$ parameter of 0.11. This is lower than the $\hat{b}$ parameters estimated in the bootstrap and bucket analysis. The lower $\hat{b}$ parameter suggests a lower sensitivity of parameterised CoVs to changes in reserve volume. It is perhaps evidence of capital modelling practitioners not

taking sufficient account of volume when selecting CoVs. A similar result was obtained when fitting a model to the 2023 benchmarking data, with a $\hat{b}$ parameter of 0.13.

Figure 17

This is further evidenced by Figure 17. The benchmark CoVs for smaller books sit below the bootstrap base model line, implying the CoVs for smaller books are low compared with those suggested by analysis of historical data. Conversely, the CoVs for very large books are slightly high compared with those suggested by analysis of historical data, and are mostly sitting above the bootstrap base model line. The 'break-even point' where the two lines intercept suggest the average capital actuary's mental reference point for volume is about $50m. This corresponds to a CoV of about 25%.

8. Results

8.1. Summary of results

The below table sets out the results from each of the analyses performed.

Method	Dataset	Model name	$\hat{b}$	R^2	# parameters	Section
Bucket analysis	**Schedule P**	**Base model**	**0.22**	**91%**	**2**	**5.1**
	Schedule P	Category model (regular)	0.22	91%	4	5.3
	Schedule P	Category model (advanced)	0.18 - 0.23	90%	6	
	Schedule P	Class model (regular)	0.22	90%	12	
	Schedule P	Class model (advanced)	0.12 - 0.28	92%	22	
Bootstrap	Schedule P	Bootstrap base model	0.29	33%	2	6.4
	Schedule P	Bootstrap class model	0.29	40%	12	
	Schedule P	Bootstrap cohort model	0.26	45%	3	
	Schedule P	Bootstrap class/cohort model	0.25	42%	13	
	Schedule P	Bootstrap bucket model	0.21	95%	2	6.5
	APRA	Bootstrap base model	0.31	66%	2	6.6
Survey data	2023 LCP Benchmarking survey	Base model	0.13	12%	2	7
	2024 LCP Benchmarking survey	Base model	0.11	12%	2	

8.2. Conclusions

Our key conclusion is that the power law model described in section 3.1 is a good fit for the relationship between volume and volatility of insurance risks, based on the analysis performed.

We recommend use of Equation 4: $CoV_{after} = CoV_{before} * \left(\frac{v_{before}}{v_{after}}\right)^b$

for adjusting volatility for changes in volume, with a b parameter of 0.22.

The reason for the selection of 0.22 is that, based on the analyses we have performed, we believe that the Schedule P dataset is sufficiently large and varied to provide credible results for general use in capital modelling, and that the bucket analysis is generally more robust than the bootstrap approach, mainly due to fact that it relies on actual reserve movements observed historically, rather than automated methods that do not take into account the specifics of each insurance portfolio considered. However, we believe consideration of the bootstrap approach is still useful in terms of understanding the relationship between volume and volatility, and in particular we are encouraged by the similarity between the Schedule P Base model and the Schedule P bootstrap bucket model.

In terms of the more complex models with more parameters, our view is that class is the main driver of differences, and that this is explained by the balance of specific and systemic risk within each portfolio. Practitioners should consider varying the value of b when carrying out modelling work, in particular if they are considering business that is materially different from the main classes within Schedule P, eg large volumes of specialty or reinsurance business. More work is needed to fully understand the impact of class on the results we have obtained.

The APRA dataset provides a useful sense check on the Schedule P results, specifically that the bootstrap base model provides similar results for Schedule P and APRA. However, the dataset is too sparse to draw meaningful conclusions on its own.

The benchmarking results indicate that capital modelling practitioners may not be taking sufficient account of volume when selecting CoVs, and may be overly anchored to volumes around the tens of millions in USD terms (equivalent to a CoV of around 25%). Volumes larger than this may therefore have CoVs that are overstated, and volumes smaller than this may have CoVs that are understated. Practitioners should be aware of this potential source of bias when parameterising, and take steps to counter it.

8.3. Applications

The below table shows some worked examples that indicate the percentage impacts of applying the above formula with the suggested b parameter of 0.22. The table shows figures for a reference starting volume of 100, and reference CoV of 30%, but the percentage impacts in the final column will be the same regardless of what volume and CoV are used.

Reference volume (v_{before})	Target volume (v_{after})	Reference CoV (CoV_{before})	Target CoV (CoV_{after})	% movement in CoV
100	10	30%	$30\% * \left({}^{10}/_{100}\right)^{-0.22} = 49.8\%$	+ 65.9%
100	25	30%	$30\% * \left({}^{25}/_{100}\right)^{-0.22} = 40.7\%$	+ 35.7%
100	50	30%	$30\% * \left({}^{50}/_{100}\right)^{-0.22} = 34.9\%$	+ 16.5%
100	75	30%	$30\% * \left({}^{75}/_{100}\right)^{-0.22} = 32.0\%$	+ 6.5%
100	125	30%	$30\% * \left({}^{125}/_{100}\right)^{-0.22} = 28.6\%$	- 4.8%
100	250	30%	$30\% * \left({}^{250}/_{100}\right)^{-0.22} = 24.5\%$	-18.3%
100	500	30%	$30\% * \left({}^{500}/_{100}\right)^{-0.22} = 21.1\%$	- 29.8%
100	1,000	30%	$30\% * \left({}^{1,000}/_{100}\right)^{-0.22} = 18.1\%$	- 39.7%

Of particular note is that, even for changes of an order of magnitude, the percentage movements in CoV are not particularly large – approximately a two-thirds increase for reducing by an order of magnitude (from 100 to 10, in this case) and approximately a 40% decrease for increasing by an order of magnitude (from 100 to 1,000, in this case). This adjustment is therefore not expected to be highly material to capital model outputs. However, we do note that, since the adjustment may need to be applied to multiple classes of business, it could form a systemic assumption within a capital modelling framework. In this case, the assumption should be tested appropriately as part of routine model validation.

Informally, we suggest that the sorts of movements shown in the table above should form the limits of the application of this approach. If volumes for a particular class of business have moved by more than an order of magnitude, it is likely that the change constitutes a change in risk profile, and therefore this method would not be appropriate.

In line with the examples noted in section 2.1, we expect that typical uses for this approach would focus on the following areas:

- Adjusting reserve risk coefficients of variation (CoVs) that were parameterised based on experience at the previous year end for future projected reserve volumes.
- Adjusting underwriting risk parameters that were calculated based on historical data, with volumes varying over the timeframe considered, but where planned future premiums are outside the range observed historically due to changes in business plans.

- Estimating appropriate volatility parameters in a systematic way for sensitivity testing that considers changes in volumes.
- Adjusting volatility parameters that have been rolled forward unadjusted for volume changes, if a full re-parameterisation is deemed unnecessary (eg if the risk profile is unchanged) or cannot be performed due to time constraints.
- Estimating volatility parameters for later years of a multi-year capital projection, or for scenarios involving changes to planned business volumes.
- Parameterising volatility distributions for new (or small) classes of business with reference to an existing (or larger) class.

This list is not intended to be exhaustive. The authors would welcome correspondence with readers who have found other uses, or who wish to discuss the methodology or implementation into a capital modelling framework.

9. Bucket analysis – uncertainties and limitations

In this section, we will focus on our reasoning for employing a bucketing technique, and the uncertainty arising from the bucketing and model fitting. There are two main sources of uncertainty arising from the bucket analysis:

- Parameter uncertainty
- Bucket uncertainty

9.1. Why do we bucket?

It is important to note that reserve volume is not the only predictor of reserve volatility. In practice, there will be many other factors at play, eg:

- Class of business
- Age of reserves
- Reserving methodology used

The process of bucketing data can be interpreted as smoothing out noise from the other driving factors of reserve volatility.

9.2. Parameter uncertainty

Parameter uncertainty is the uncertainty that arises in estimating $\hat{b}$ using the linear model. It arises because the value $b = 0.22$ obtained is a statistical estimate based on the bucketed data. We can directly measure this source of uncertainty as follows:

The following formula is used to calculate the residual sum of squares:

$$SS_{res} = \sum_j (\sigma_j - \hat{\sigma}_j)^2$$

where σ_j is the CoV, the response variable and $\hat{\sigma}_j$ is the predicted CoV for each bucket B_j.

We can then estimate the variance of the residuals $\hat{\tau}^2$ using the following formula:

$$\hat{\tau}^2 = \frac{SS_{res}}{n_\Sigma - 2}$$

where n_Σ is the number of data points, in this case, the number of buckets B_j.

Then we have:

$$SE(\hat{b}) = \sqrt{\frac{\hat{\tau}^2}{\sum_j (x_j - \bar{x})^2}} = \frac{1}{\sqrt{n_\Sigma - 2}} \sqrt{\frac{\sum_j (y_j - \bar{y})^2}{\sum_j (x_j - \bar{x})^2}}$$

Equation 11

where x_j is the predictor variable, which in this case is $\log v_j$, ie the logarithm of the mean reserve volume for each bucket B_j.

Calculating this for the base model, we obtain $SE(\hat{b}) = 0.004$.

Given that $SE(\hat{b}) = 0.004 \ll 0.22 = \hat{b}$, we are comfortable that this source of error has a low materiality on the results of the analysis.

9.3. Bucket uncertainty

While parameter uncertainty is widely known and well understood, the bucket analysis has a second, more nuanced source of uncertainty. In the process of fitting a linear model to the bucketed data points, we are implicitly assuming that the positions of the bucketed data points are correct.

Recall within each bucket B_j, we calculated the mean reserve size $\bar{R}_j$ and the volatility σ_j of the reserve movements within the bucket using the following formulae:

$$\bar{R}_j = \frac{1}{n_B}\sum_{B_j} R_i$$

$$\bar{r}_j = \frac{1}{n_B}\sum_{B_j} r_i$$

$$\sigma_j^2 = \frac{1}{n_B - 1}\sum_{B_j}\left(r_i - \bar{r}_j\right)^2$$

Where r_i is the reserve movement for datapoint (R_i, r_i) within bucket B_j.

We plot the volatility σ_j of the reserve movements against the mean reserve size $\bar{R}_j$, and fit a line equivalent to:

$$\sigma = a\bar{R}^{-b}$$

Which is equivalent to Equation 1:

$$CoV = av^{-b}$$

where σ is the CoV and $\bar{R}$ is the volume v.

We can therefore consider mean reserve size $\bar{R}_j$ to be a 'horizontal' coordinate and volatility σ_j a 'vertical' coordinate.

Parameter uncertainty measures the statistical uncertainty associated with the estimated value $\hat{b}$. However, this does not allow for the fact that we have made statistical estimates in the calculation of the horizontal and vertical coordinates of each bucketed data point $(\bar{R}_j, \sigma_j)$.

We have therefore defined bucket uncertainty as the uncertainty that arises in quantifying the reserve size $\bar{R}_j$ and volatility σ_j for each bucket B_j. We can informally think of this as the uncertainty in the coordinates of each bucketed data point.

9.4. Vertical bucket uncertainty

To measure the uncertainty in the vertical coordinate of each bucketed data point, we can measure the standard error of the volatility estimate.

Note that we can calculate the standard error of the variance estimate σ_j for each bucket B_j as:

$$SE(\sigma_j) = \frac{\sigma_j}{\sqrt{2n_B}} * \sqrt{1 - \frac{1}{n_B}}$$

It therefore follows that, if we measure the error term $SE(\sigma_j)$ as a proportion of the volatility σ_j, we get:

$$\frac{SE(\sigma_j)}{\sigma_j} = \frac{1}{\sqrt{2n_B}} * \sqrt{1 - \frac{1}{n_B}}$$

Which, for a constant bucket size of $n_B = 200$, leads to:

$$\frac{SE(\sigma_j)}{\sigma_j} = \frac{1}{\sqrt{2n_B}} * \sqrt{1 - \frac{1}{n_B}} = 5\%$$

which is independent of the volatility σ_j.

Referring back to Figure 8, we can define vertical error bars for each bucketed data point representing the standard error $SE(\sigma_j)$ above and below the bucket volatility σ_j, such that on a logarithmic scale, the error bar for bucket B_j has length:

$$\log\left(\sigma_j + SE(\sigma_j)\right) - \log\left(\sigma_j - SE(\sigma_j)\right)$$

We can visualise the vertical bucket uncertainty by including these vertical error bars on Figure 8 as shown below:

Figure 18

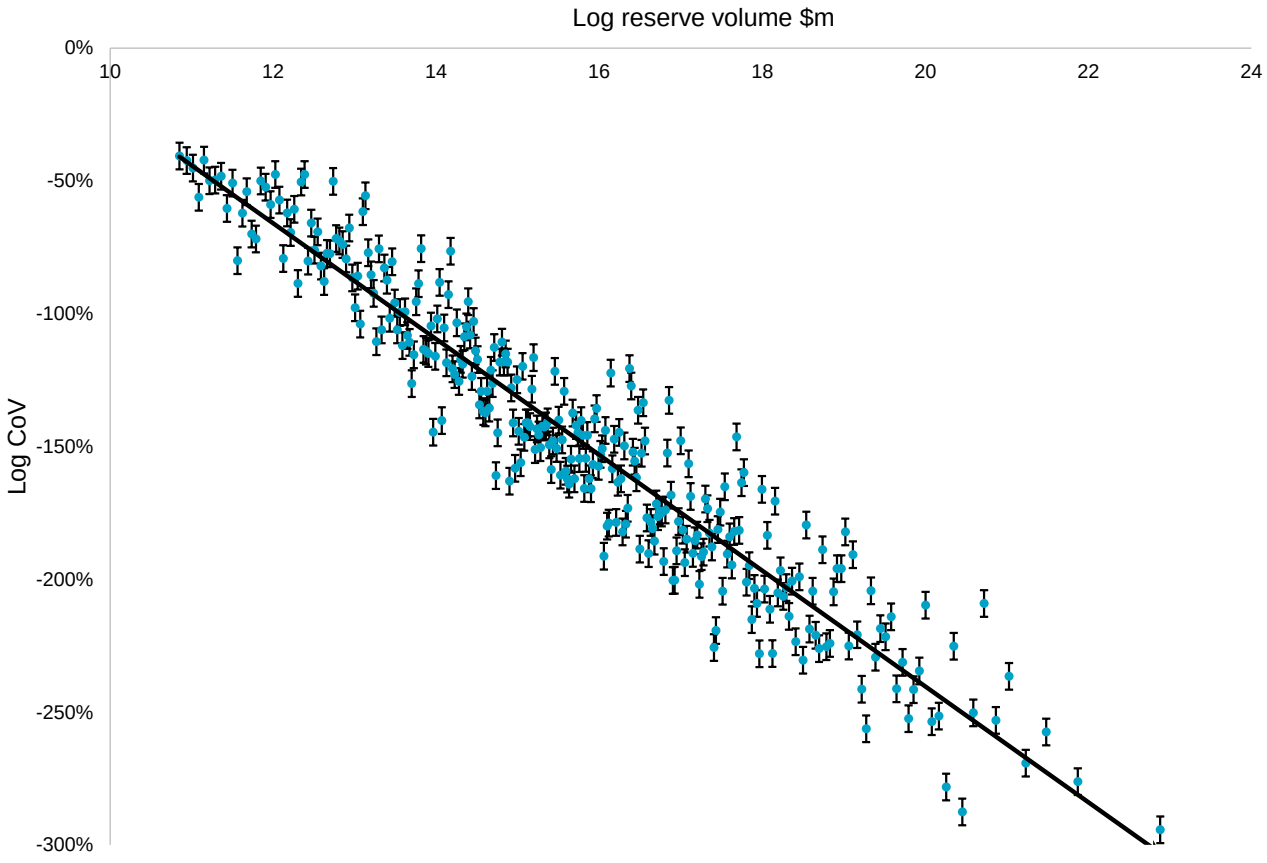

Note that the vertical error bars are the same size for each bucketed data point.

We can understand this phenomenon mathematically. The total size of the vertical error bar for a given bucket B_j is given as:

$$\log\left(\sigma_j + SE(\sigma_j)\right) - \log\left(\sigma_j - SE(\sigma_j)\right)$$

$$= \log(1.05 * \sigma_j) - \log(0.95 * \sigma_j)$$

$$= \log\left(\frac{1.05 * \sigma_j}{0.95 * \sigma_j}\right) = \log\left(\frac{1.05}{0.95}\right) = 0.100$$

which is independent of B_j, and therefore the size of the vertical error bar is constant across all bucketed data points.

The gradient of the slope is affected by this vertical coordinate, and where the vertical coordinate volatility σ_j sits within the estimate. It is possible that, if the volatility σ_j was estimated at a different point along the range of the standard error, the slope would shift and therefore the estimated b parameter from the fitted model would change. Therefore, additional validation is required.

We performed a stochastic analysis to understand how stable the value of b is to perturbations in the vertical coordinate of the bucketed data points. In this way, we can quantify vertical bucket uncertainty for a given bucket size, using the analysis performed above to quantify the uncertainty in the vertical coordinate of each bucketed data point.

For each bucketed data point $(\bar{R}_j, \sigma_j)$, we transformed the vertical coordinate into a stochastic random variable $(\bar{R}_j, \Theta_j)$ using the following approach:

$$\Theta_j \sim N\left(\sigma_j, SE(\sigma_j)^2\right)$$

The following analysis is a justification of the choice of a normal distribution for the stochastic analysis:

Treating reserve movements r_i as random data points, we have that

$$\sigma_j^2 = \frac{1}{n_B - 1}\sum_{i \in B_j}\left(r_i - \bar{r}_j\right)^2 \sim k\chi^2_{n_B - 1}$$

Ie our volatility metric for each bucket has a chi-squared distribution with $n_B - 1$ degrees of freedom, multiplied by some constant k. Then note that we can use the following result:

If $X \sim \chi^2_n$ then $\sqrt{2X}$ is approximately normally distributed with mean $\sqrt{2n-1}$ and variance 1.[3]

Therefore we can approximate Θ_j with a normal distribution, with mean σ_j and standard deviation given by $SE(\sigma_j)$.

We then used the same approach of fitting a linear model to the log-transformed data $\{(\bar{R}_j, \Theta_j)\}_j$ and measured the slope coefficient to find $\hat{b}$.

We repeated this for 1,000 simulations using a Monte Carlo simulation approach. For each simulation, we obtained a value for $\hat{b}$.

[3] R.A Fisher, 1922

Figure 19 shows the empirical CDF of the $\hat{b}$ values obtained:

Figure 19

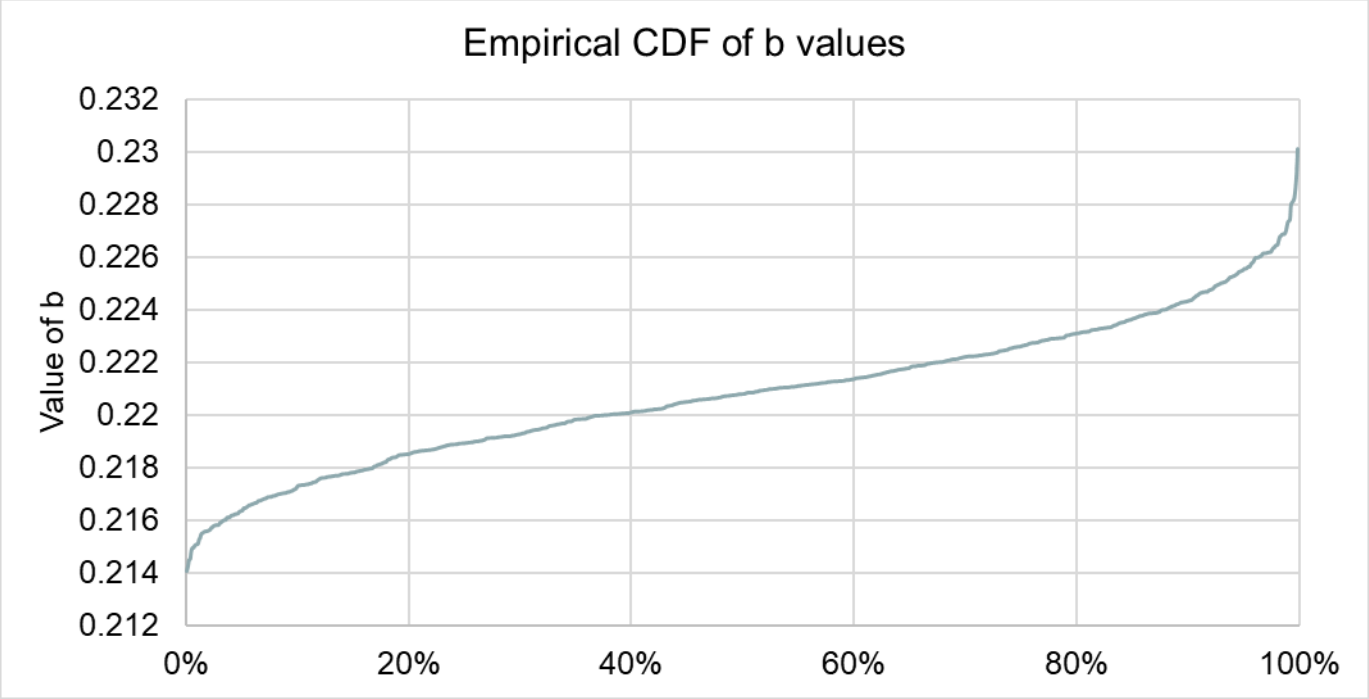

We can see that there is a tight spread of $\hat{b}$ estimates around 0.22. The sample standard deviation of the $\hat{b}$ estimates is 0.0026, which as a percentage of our original $\hat{b}$ estimate is

$$\frac{0.0026}{0.22} \approx 1\%$$

Given that this is low, we are comfortable that this source of error has a low materiality on the results of the analysis.

9.5. Horizontal bucket uncertainty

To measure the uncertainty in the horizontal coordinate of the bucketed data points, we can measure the standard error of the volume estimate.

We can calculate the standard error of the sample mean estimate for each bucket B_j as

$$SE(\bar{R}_j) = \frac{1}{\sqrt{n_B}} \sqrt{\frac{1}{n_B - 1} \sum_{B_j} (R_i - \bar{R}_j)^2}$$

where R_i represents the actual reserve volume for data point i within bucket B_j.

In a similar way to vertical bucket uncertainty, we can define horizontal error bars for each bucketed data point representing the standard error $SE(\bar{R}_j)$ above and below the bucket reserve volume $\bar{R}_j$, such that on a logarithmic scale, the error bar for bucket B_j has length

$$\log\left(\bar{R}_j + SE(\bar{R}_J)\right) - \log\left(\bar{R}_j - SE(\bar{R}_j)\right)$$

Note that, unlike the normalised standard error terms for vertical bucket uncertainty, the value of these normalised standard error terms for the horizontal bucket uncertainty varies across the buckets. We can plot the length of these error bars against the corresponding log mean reserve size $\log \bar{R}_j$ to obtain the following chart:

Figure 20

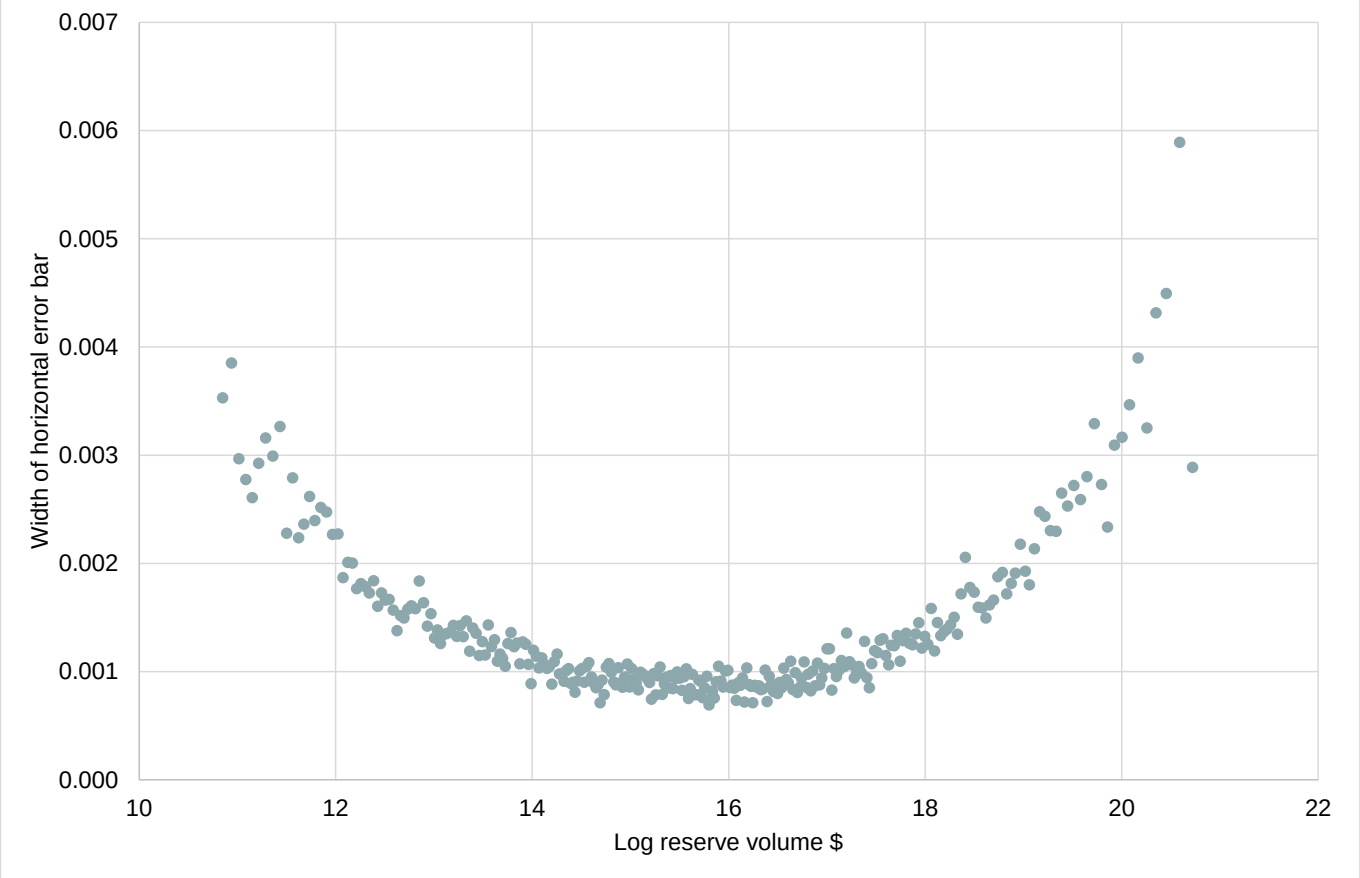

There is a clear 'U' shape in the above chart. We can see that the uncertainty in the horizontal coordinate of bucketed data points is the highest for the buckets corresponding to the smallest and largest reserve volumes. There is higher uncertainty for the largest reserve volumes because, as previously mentioned, there is a large spread of reserve volumes within these buckets due to the sparsity of data for these high reserve volumes. There is also higher uncertainty for the smallest reserve volumes, as the standard error of $\bar{R}_j$ for these buckets is high as a proportion of the low reserve volume. The graph shows that 0.006 is an upper bound for the horizontal error bar widths. This means that the horizontal error bars are significantly narrower than the vertical error bars ($0.006 \ll 0.1$).

We can also consider a second form of horizontal uncertainty. As part of the bucketing process, recall that we calculated the reserve volume in each bucket using the following formula:

$$\bar{R}_j = \frac{1}{n_B} \sum_{B_j} R_i$$

The choice of taking the arithmetic mean as the representative reserve volume for each bucket was in some ways arbitrary. Alternatively, we could have chosen a number of other metrics, including:

- Minimum
- Maximum
- Median
- Geometric mean

Using one of the above alternative metrics would result in the horizontal coordinate of each bucketed data point lying at a different point in the bucket. Therefore, we can quantify the secondary source of horizontal bucket uncertainty as the maximum width of each bucket:

$$Bucket\ width(B_j) = \max_{B_j} R_i - \min_{B_j} R_i$$

Equation 12

Note that the magnitude of this uncertainty will increase with bucket size. The original dataset Γ is such that most of the data-points (R_i, r_i) are of smaller reserve volumes R_i, as shown in the histogram below:

Figure 21

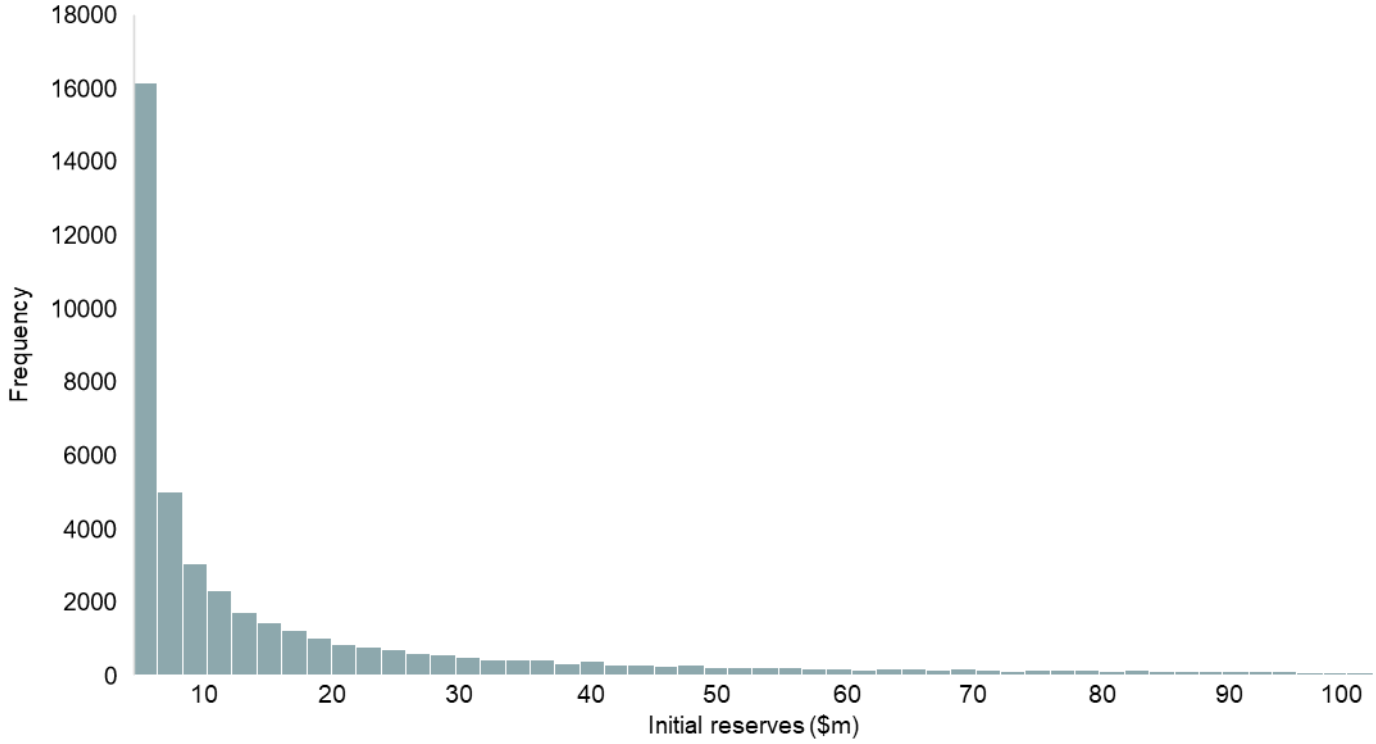

Recall that we have fixed the size of each bucket to contain exactly 200 data points, ie $n_B = 200$. Therefore, the sparsity of data for higher reserve volumes results in wider buckets at these volumes.

We can redefine the horizontal error bars mentioned above such that for each bucket, the width of the error bar is given by the bucket width in Equation 12.

Then in a similar way to the vertical bucket uncertainty, we can plot the horizontal error bars associated with the secondary horizontal bucket uncertainty:

Figure 22

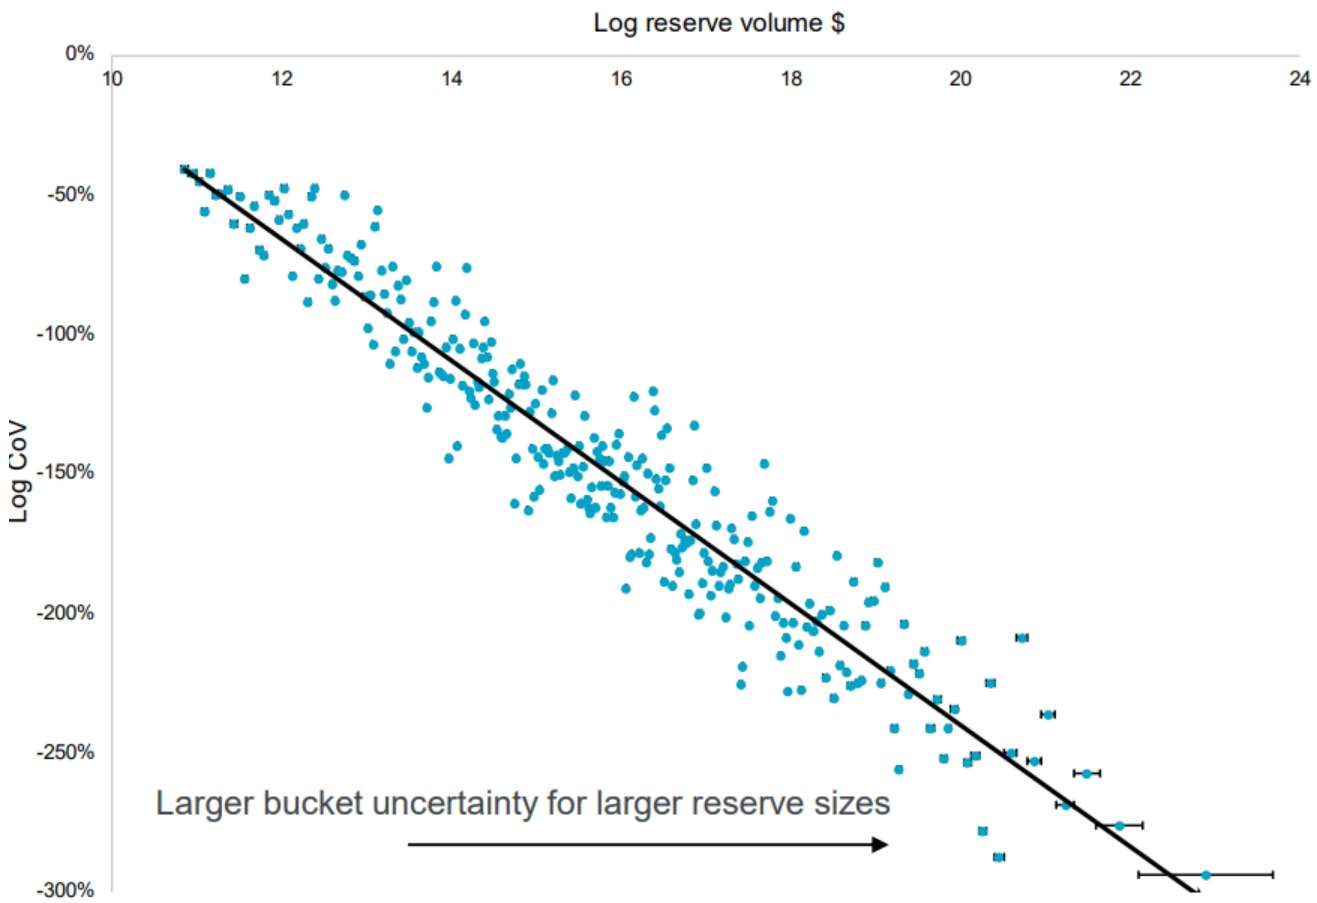

Visually, we can see that the uncertainty bars are almost negligible for the vast majority of the bucketed data points. It is only for the bucketed data points with the largest reserve volumes where we can see the uncertainty bars.

We can perform additional validation by analysing the stability of the calculated b parameter when using the alternative metrics noted above.

For each alternative reserve volume metric stated above, we created a new dataset from which to fit a model:

- $\Sigma_{minimum}$
- $\Sigma_{maximum}$
- Σ_{median}
- $\Sigma_{geometric\ mean}$

This was done by calculating a representative volume $\bar{R}_j$ for each bucket B_j via the specified metric.

Using the minimum metric as an example, to create $\Sigma_{minimum}$, for each reserve volume R_i in bucket B_j, the representative reserve volume for the bucket is:

$$\bar{R}_{j_{min}} = Min_{i \in B_j}\{R_i\}$$

Calculating the volatility σ_j as before, we can obtain the dataset:

$$\Sigma_{\text{minimum}} = \{(\bar{R}_{j_{min}}, \sigma_j)\}_j$$

from which we can fit a model to parameterise b as described previously.

The different b parameters obtained from using the different reserve volume metrics are shown in Figure 23.

Figure 23

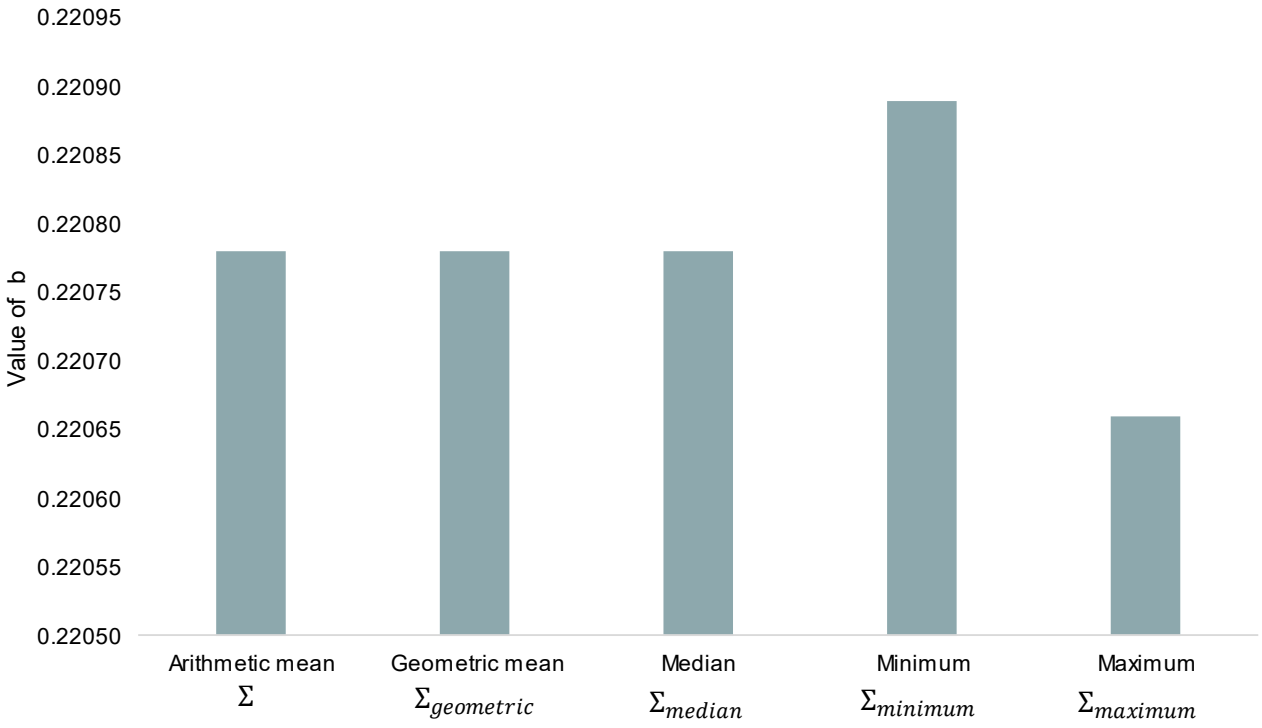

We can see that the estimate of $\hat{b}$ is very stable to calculate the reserve volume for each bucket under different metrics, as the largest deviation has a $\hat{b}$ value difference of 0.00023 (around 0.1% of the estimate for b).

9.6. Selection of bucket size n_B

In choosing the size of the buckets, it is important to consider how the choice of bucket size impacts the magnitude of the different sources of uncertainty highlighted in the above sections. We can note the following dynamics when choosing bucket size:

Recall in section 9.4, we quantified the uncertainty in the volatility estimate σ_j for each bucket as a proportion of σ_j as follows:

$$\frac{SE(\sigma_j)}{\sigma_j} = \frac{1}{\sqrt{2n_B}} * \sqrt{1 - \frac{1}{n_B}}$$

Note that this is a decreasing function in n_B for $n_B \geq 2$, and therefore as bucket size increases, the uncertainty in the response variable of the linear model σ_j decreases.

We can take this one step further by analysing the relationship between overall bucket uncertainty and the chosen bucket size.

We can define overall bucket uncertainty using a similar approach as taken in section 9.4, however this time also including primary horizontal bucket uncertainty:

For each bucketed data point $(\bar{R}_j, \sigma_j)$, we can transform the vertical coordinate into a stochastic random variable (Δ_j, Θ_j) using the following approach:

$$\Delta_j \sim N\left(\bar{R}_j, SE\left(\bar{R}_j\right)\right) \qquad \text{and} \qquad \Theta_j \sim N\left(\sigma_j, SE\left(\sigma_j\right)\right)$$

Again, we can justify the choice of a normal distribution for Θ_j in the same way as set out in section 9.4.

Note that by the central limit theorem, treating $R_i \in \Gamma$ as random data points, we can approximate $\bar{R}_j$ with a normal distribution.

We then used the same approach of fitting a linear model to the log-transformed data $\{(\Delta_j, \Theta_j)\}_j$ and measured the slope coefficient to find $\hat{b}$.

We repeated this for 1000 simulations using a Monte Carlo simulation approach. For each simulation, we obtained a value for $\hat{b}$, and therefore for a given bucket size, we create an empirical distribution of $\hat{b}$ parameters. We define overall bucket uncertainty as the standard deviation of this empirical distribution.

We can repeat this process across different bucket sizes and obtain values for overall bucket uncertainty for each bucket size. The chart below shows the relationship between overall bucket uncertainty and bucket size:

Figure 24

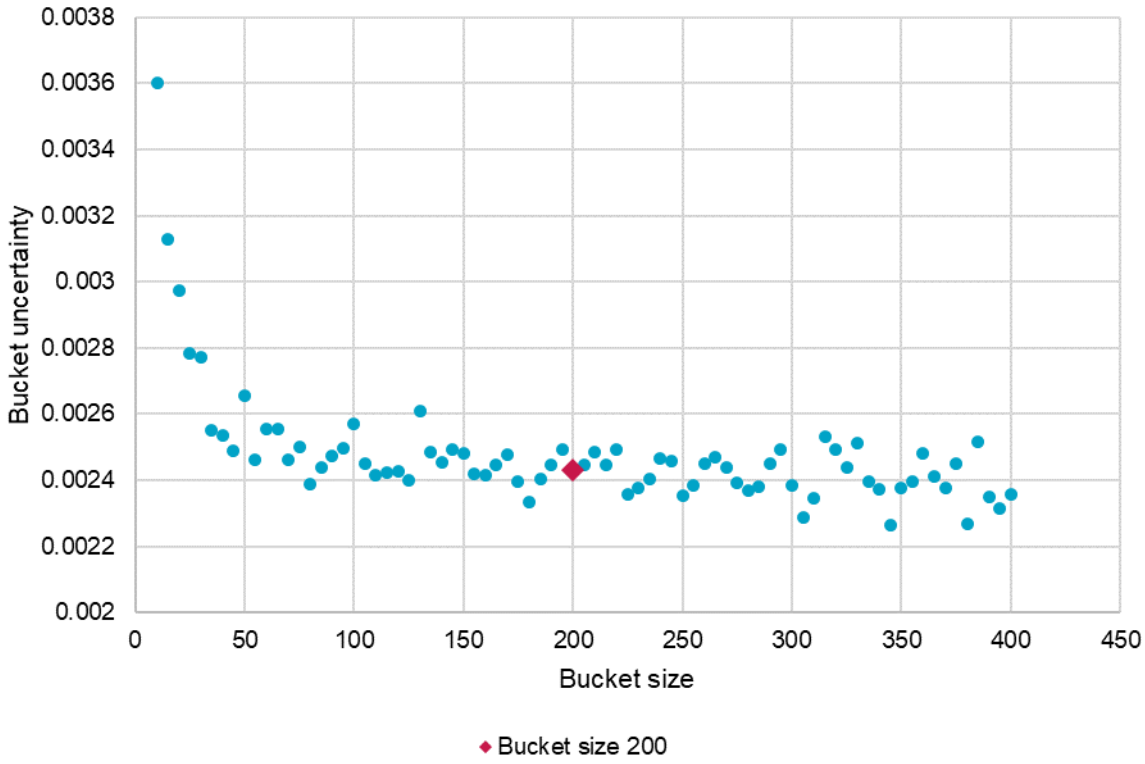

We can see that there is a large improvement in bucket uncertainty as bucket size increases from 1 to 50. After this point, there is no clear evidence that increasing bucket size materially improves the level of bucket uncertainty in the analysis.

We fit models to these bucketed datasets to examine this further:

- We performed the bucketing analysis as described in section 5.1 on a range of bucket sizes from 5 to 450 in steps of 5.
- In this way, we obtained bucketed data sets $\Sigma_5, \Sigma_{10}, \ldots, \Sigma_{450}$, where Σ_n denotes the bucketed data set corresponding to a bucket size of n.

- For each bucketed data set, we fit a model using linear regression in the same way as described in section 5.1.
- We extracted the parameter estimate of $\hat{b}$ and the parameter uncertainty $SE(\hat{b})$ in each case.

Figure 25 shows the effect of bucket size on parameter uncertainty of $\hat{b}$:

Figure 25

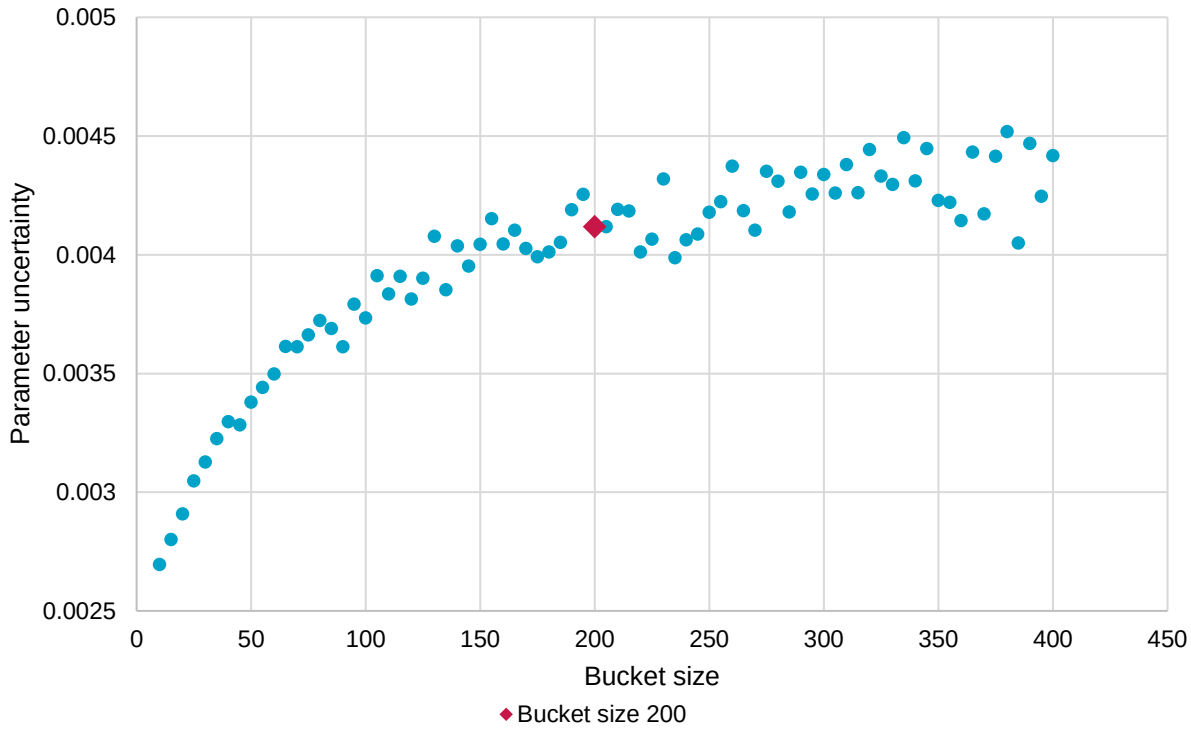

Bucket size $n_B = 200$ used in our analysis is highlighted on the chart.

From this graph, we can see that parameter uncertainty increases as bucket size increases.

Using the fitted models, we looked at the relationship between bucket size and the fitted $\hat{b}$ parameter:

Figure 26

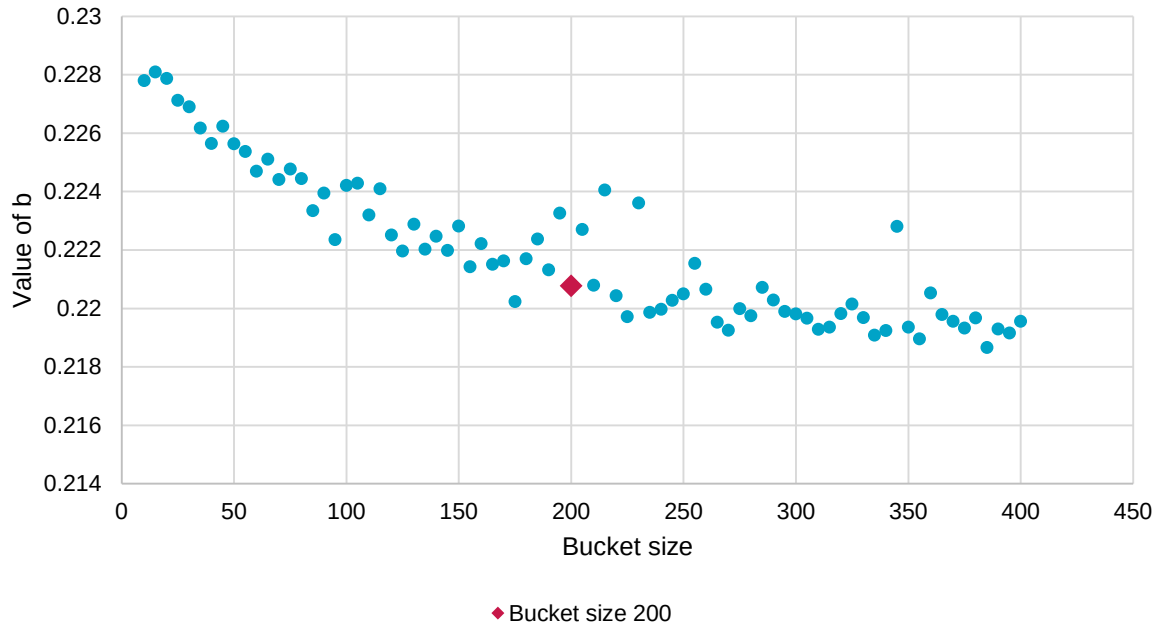

Figure 26 shows how the derived b parameter changes across different choices of bucket size. We can observe that as bucket size increases, the derived b parameter tends to decrease. While it is not immediately obvious or intuitive why this trend exists, we note that the effect is small: all bucket sizes considered greater than around 75 result in b parameters of 0.22 when rounding to two decimal places, and the trend appears to stop (or at least slow) for bucket sizes greater than around 200.

Finally, we investigated the trade-off between overall bucket uncertainty and parameter uncertainty:

Figure 27

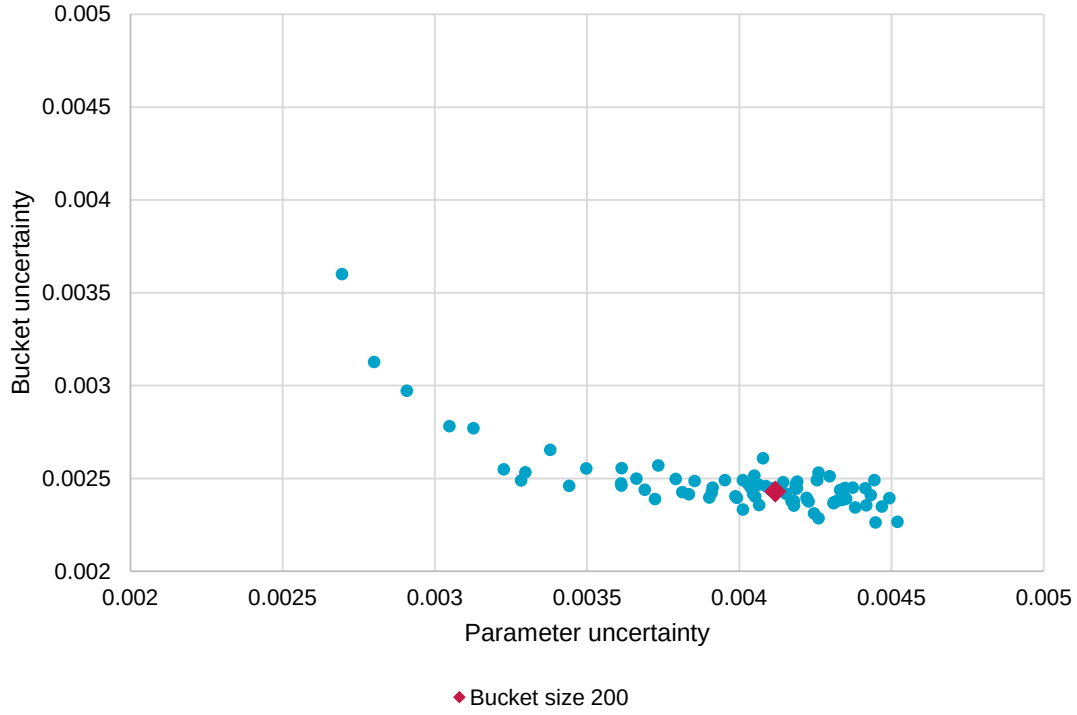

Figure 27 highlights the trade-off between bucket uncertainty and parameter uncertainty. Each point on the graph corresponds to a bucket size. The bucket sizes plotted range from 15 to 450. From the graph, we can see that there is a clear trade-off between the two uncertainties for large bucket uncertainty and small parameter uncertainty; however, the trend does not continue for higher parameter uncertainties. This demonstrates that increasing bucket size beyond a certain point will increase parameter uncertainty without providing the benefit of reduced bucket uncertainty. In this way, we can see that there is a natural upper bound for the optimal bucket size. The point on the graph corresponding to a bucket size of 200 is highlighted to show where our chosen bucket size sits on the trade-off curve. We can see that this bucket size provides a very low bucket uncertainty, and a reasonably low parameter uncertainty. While it is possible to further reduce parameter uncertainty without being penalised too heavily on bucket uncertainty, the reduction in parameter uncertainty has very little impact on the results of the analysis.

9.7. Conclusion

Combining the results from above, we have the trade-off that as parameter uncertainty increases, bucket uncertainty decreases. Intuitively this makes sense, as we would expect that, as you increase bucket size, the number of data points within each bucket increases and therefore the statistical estimates of volume and volatility of each bucket become less uncertain. On the other hand, increasing the bucket size results in a lower number of buckets to fit a linear model to, and therefore the model parameter estimates are more uncertain.

Appendices

Appendix 1: R-squared

The R-squared value for a fitted linear model can be defined using the following equations:

$$R^2 = 1 - \frac{SS_{res}}{SS_{tot}}$$

$$SS_{res} = \sum_i (y_i - \hat{y}_i)^2$$

$$SS_{tot} = \sum_i (y_i - \bar{y})^2$$

Where y_i is the response variable, $\hat{y}_i$ is the predicted value and $\bar{y} = \frac{1}{n} * \sum y_i$

In our case, we have

$$y_i = \log CoV_i$$

$$\hat{y}_i = \hat{a} - \hat{b} * \log v_i$$

$$\bar{y} = \frac{1}{n_\Sigma} * \sum_i \log CoV_i$$

The R^2 value is defined as the proportion of the variation in the dependent variable that is predictable from the independent variables. Therefore, a higher value of R^2 indicates a better model fit and we have the following numerical bounds for R^2:

$$0 < R^2 < 1$$

Appendix 2: Over-dispersed Poisson distribution framework

Unlike the Poisson distribution, where the variance equals the mean, the over-dispersed Poisson distribution has a variance proportional to the mean. In claims reserving, this model assumes that the incremental claims C_{ij} are independent over-dispersed Poisson random variables with mean and variance:

$$E[C_{ij}] = \mu_{ij} = x_i y_j \text{ and } Var[C_{ij}] = \phi\, x_i y_j$$

Where x_i is the expected ultimate claims and y_j is the proportion of ultimate claims emerging in each development period. The ϕ parameter introduces the over-dispersion which is estimated from the data. Further details of this can be found in Renshaw & Verrall (1998)[4] and in England & Verrall (2002)[5].

The main limitation of the bootstrap is that the sum of the incremental claims for each development period must be positive. This is due to the y_j in the variance parameter. Some negative increments can occur as long as the sum of any column is not negative.

The log link function in the GLM is introduced to re-parameterise the mean structure so that it has a linear form such that:

$$log(\mu_{ij}) = c + \alpha_i + \beta_j$$

Since there is a parameter for each row i and column j, the structure is still a chain ladder type.

[4] *Renshaw, A. E. & Verrall, R. J. (1998). A stochastic model underlying the chain-ladder technique. B.A.J. 4, 903-923*
[5] *England, P.D. & Verrall, R.J. (2002). Stochastic claims reserving in general insurance. Institute of Actuaries*

Appendix 3: Bootstrapping methodology

Fitting the GLM

Given an incremental claims triangle, we can calculate an expected triangle, from which Pearson residuals can also be calculated. We define η_{ij} to be the loglink of μ_{ij}, where the loglink function is defined as:

- $\ln(\mu_{ij})$ for $C_{ij} > 0$
- 0 for $\mu_{ij} = 0$
- $-\ln(-\mu_{ij})$ for $C_{ij} < 0$

We then further define $\eta_{ij} = \alpha_i + \beta_j$, where $i = 1, 2, \dots, n$ and $j = 1, 2, \dots, n$, and define a system of equations as $Y = XA$ to calculate η_{ij}, with:

$$Y = \begin{pmatrix} \ln(C_{11}) \\ \ln(C_{21}) \\ \ln(C_{31}) \\ \vdots \\ \ln(C_{n1}) \end{pmatrix}$$

$$X = \begin{pmatrix} 1 & 0 & 0 & \cdots & 0 \\ 0 & 1 & 0 & \cdots & 0 \\ 0 & 0 & 1 & \cdots & 0 \\ \vdots & \vdots & 0 & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & 1 \end{pmatrix}$$

$$A = \begin{pmatrix} \alpha_1 \\ \alpha_2 \\ \alpha_3 \\ \vdots \\ \alpha_n \end{pmatrix}$$

We can solve this using a least squares approach for the chosen distribution, and the calculate the fitted matrix, $\hat{Y} = XA$, which will be of the form:

$$\hat{Y} = \begin{pmatrix} \alpha_1 \\ \alpha_2 \\ \alpha_3 \\ \vdots \\ \alpha_1 + \beta_2 + \cdots + \beta_n \end{pmatrix} = \begin{pmatrix} \eta_{11} \\ \eta_{21} \\ \eta_{31} \\ \vdots \\ \eta_{n1} \end{pmatrix}$$

This can then be exponentiated using the inverse of the loglink function to find u_{ij}, the expected values for each entry of the triangle according to the GLM as:

- $\exp(\eta_{ij})$ for $C_{ij} > 0$
- 0 for $C_{ij} = 0$
- $-\exp(\eta_{ij})$ for $C_{ij} < 0$

For a more comprehensive explanation on this approach, refer to Shapland, 2016[6].

Residuals

[6] *Shapland, M., 2016. Using the ODP Bootstrap Model: A Practitioner's Guide. Casualty Actuarial Society*

After calculating all values of μ_{ij}, we can now calculate residuals for all entries in the incremental triangle, from which sampling can begin. Firstly, the unscaled Pearson residuals can be calculated as:

$$r_{ij} = \frac{C_{ij} - \mu_{ij}}{\sqrt{|\mu_{ij}^z|}} * f^{DoF}$$

f^{DoF} is an adjustment made to effectively allow for over-dispersion of the residuals in the sampling process and to add process variance to approximate a distribution of possible outcomes. It is defined as $f^{DoF} = \frac{N}{N-P}$, where N and p are defined below.

The scale parameter is calculated as either of the following two formulae, depending on whether it is set to be constant by the user, or changing for each development period j:

$$\phi = \sqrt{\frac{\sum_{ij} r_{ij}^2}{N - p}}$$

or

$$\phi_j = \sqrt{\frac{\sum_i r_{ij}^2}{o_j}},$$

where N is the number of observations or incremental data cells in the triangle, p is the number of parameters and o_j is the number of entries in the incremental triangle for that development period. Typically, $N = \frac{n(n+1)}{2}, p = 2(n-1)$ and $o_j = n + 1 - j$ for an $n \times n$ triangle.

Sampling residuals

The residuals are then randomly sampled from with replacement and used to create a pseudo triangle, using the following formula:

$$C_{ij}^* = r_{ij}^* * \phi_j * \sqrt{|\mu_{ij}^z|} + \mu_{ij}$$

where C_{ij}^* is our resampled residual, and r_{ij}^* is the new incremental value for the pseudo triangle and origin period i and development period j.

The user can decide if negative incremental values are allowed in the bootstrapped triangle. If they are not allowed, all negative incremental values will be set to 0. The user could also decide to prevent negative cumulative values in the bootstrapped triangle, even if negative incremental values are allowed.

Refitting the GLM

The pseudo triangle is then projected across all future development periods using the basic chain ladder approach in order to calculate a completed triangle. The link ratios are calculated across all origin years in the pseudo triangle and this is used to derive an estimate for the ultimate claims in the pseudo triangle and a reserve for each origin year can be calculated by subtracting the latest paid in the pseudo triangle from the calculated ultimate claims.

Ultimate

From the set of reserves calculated using the bootstrap method the standard deviation, mean, and CoV can be calculated.

Appendix 4: Performance of theoretical model on real world data

Using the same first principles analysis as discussed in section 3.2, with additional assumptions in relation to the correlation between claims, we derived an implied relationship between the CoV of a theoretical portfolio of claims and the corresponding mean volume of the portfolio. In this section, we discuss the derivation of this theoretical relationship based on probabilistic first principles and how the relationship compares to the 'real world' Schedule P data.

Approach

Recall from Section 3.2, we defined X_i to be the amount of the i-th claim.

Again, assume X_i are identically distributed with mean μ and variance σ^2.

Then define S as the total reserves in the following way:

$$S = \sum_{i=1}^{N} X_i$$

Given that each claim is identically distributed, we then defined the volume as total reserves to be: $v = E[S] = n\mu$.

For the derivation of the theoretical model, we now assume a constant level of correlation between claims:

$$Corr(X_i, X_j) = \rho$$

We still have:

$$E[S] = E\big[E[S|N]\big]$$

$$= E[N] * E[X_i]$$

$$= n\mu$$

But now we also have

$$Var(S) = E[Var(S|N)] + Var(E[S|N])$$

$$= E\left[Var\left(\sum_{i=1}^{N} X_i \,|N\right)\right] + Var(E[\textstyle\sum_{i=1}^{N} X_i \,|N])$$

$$= E\left[\sum_{i=1}^{N} Var(X_i) + 2\sum_{i\neq j}^{N} \sqrt{Var(X_i) * Var(X_j)} * Corr(X_i, X_j)\right] + Var(N * E[X_i])$$

Note that there are $\binom{N}{2} = \frac{N(N-1)}{2}$ such pairs $(i, j : 1 \leq i \leq N, 1 \leq j \leq N, i \neq j)$, therefore:

$$Var(S) = E[N * Var(X_i)] + E[N(N-1)\sigma^2\rho] + E[X_i]^2 * Var(N)$$

$$= n\sigma^2 + 2\sigma^2\rho * (E[N^2] - E[N]) + n\mu^2$$

$$= n\sigma^2 + \sigma^2\rho * n^2 + n\mu^2$$

$$= n(\sigma^2 + \mu^2) + n^2\rho\sigma^2$$

Allowing the calculation of CoV of the reserves as:

$$CoV(S_n) = \frac{\sqrt{Var(S)}}{E[S]} = \frac{\sqrt{n(\sigma^2 + \mu^2) + n^2\rho\sigma^2}}{n\mu}$$

$$= \frac{1}{\mu}\sqrt{\rho\sigma^2 + \frac{\sigma^2 + \mu^2}{n}}$$

$$= \frac{1}{\mu}\sqrt{\rho\sigma^2 + \frac{\sigma^2\mu + \mu^3}{n\mu}}$$

Simplifying this further, since μ, σ and ρ are the fixed parameters associated with the distribution of claims X_i, we obtain

$$CoV = \sqrt{c + \frac{d}{v}}$$

For constants $c = \frac{\sigma^2}{\mu^2}\rho$ and $d = \frac{\sigma^2}{\mu} + \mu$

Fitting this model

We can find parameters c, d to fit this model to the bucketed data points in Σ, and as such, treat this as an alternative model to our original model:

$$CoV = av^{-b}$$

Given that log-transforming the data for the alternative model no longer results in a linear relationship, the original method of fitting a linear model no longer works.

Instead, we set up an optimisation problem to coefficients a, b to minimize

$$MSE = \sum_j \left(g(\bar{R}_j) - \tilde{\sigma}_j\right)^2$$

Where the function g is defined as

$$g: x \to \sqrt{c + \frac{d}{x}}$$

Using the gradient descent algorithm to solve this optimisation problem results in parameters

$$c = 0.037$$

$$d = 39400$$

We can repeat the same optimisation problem on our original model given in Equation 1, defined by function h:

$$h: x \to ax^{-b}$$

Again using the gradient descent algorithm, the minimum found by the algorithm results in parameters

$$a = 7.63$$

$$b = 0.22$$

Given that these a and b estimates are the same as those found by fitting a linear model, the result gives comfort that this approach results in reasonable parameter estimates.

Note that the MSE for the alternative model is 0.727, which is higher than the corresponding MSE of 0.14 for the base model.

We can plot this alternative model in logarithmic space to show how closely it fits the bucketed data points.

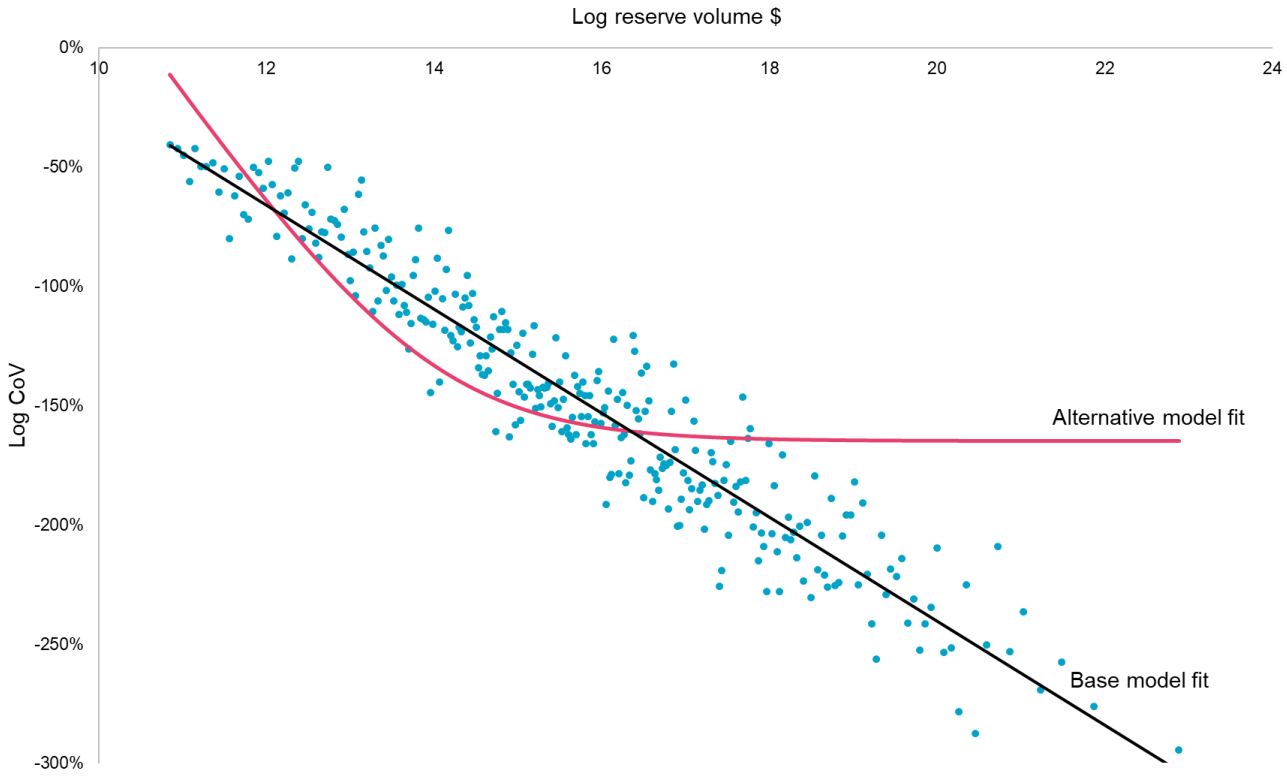

Visually, we can see that the model fit is clearly worse than the original model, particularly as reserve volume gets very large. It is therefore interesting to note that the model derived from theoretical first principles does not capture the structure of the real-world data. The discrepancy between the theory and real-world data is characterised by the horizontal levelling off of the alternative model trend line for larger reserve volumes. We can infer that the theory implies that there should be a degree of undiversifiable volatility whereas the real-world data provides no evidence that undiversifiable volatility exists.

We can calculate the level of undiversifiable volatility implied by the alternative model:

Recall we have

$$CoV = \sqrt{c + \frac{d}{V}}$$

So then

$$\lim_{V \to \infty} CoV = \lim_{V \to \infty} \sqrt{c + \frac{d}{V}}$$

$$= \sqrt{c} = \sqrt{0.037} = 19.2\%$$