British Actuarial Journal (2025), Vol. **30**, e35, pp. 1–13
doi:10.1017/S1357321725100196

Institute and Faculty of Actuaries

SESSIONAL MEETING DISCUSSION

Practical considerations in claims inflation estimation

[Institute and Faculty of Actuaries, Sessional Webinar, Monday 17 March 2025]

Mr C. T. Creedon, F.I.A.: Thank you to everyone for joining us today to discuss "Practical Considerations in Claims Inflation Estimation." The Working Party originally presented this work at GIRO in 2023 and the paper underpinning the work was runner-up in the IFoA's Brian Hey Prize in 2024. There are five of us on the Claims Inflation Working Party. Not present today are Marcus Schofield and Richard Stock. I am the chair, Cian Creedon. I am an actuary at AON, the marine broking team.

Miss E. Bargate, F.I.A.: Hi. I am Erin (Bargate). I work at Hiscox, where I am a senior underwriter in the alternative risk team.

Mr S. Lenney, F.I.A.: Hi, everyone. My name is Shane Lenney and I work at Scor in the portfolio management team.

Mr Lenney: This working party came together to try to create a pragmatic approach to estimating claims inflation, considering the availability of data and current approaches, and to see how these elements can fit together.

In today's talk, we will cover the data, Scenarios that we used to do our testing, and the evaluation of our methods and findings as we were doing these studies.

To understand the approaches, outputs and limitations, we wanted to create a set of data that we could use. We could set various inputs to see how well a range of methods could be used to produce the results; or, conversely, which ones failed or were less likely to give us a consistent result.

To create this data, we put together a tool in Excel. It is a stochastic tool that generates losses. The Visual Basic for Applications (VBA) code we created is relatively straightforward to translate into R or Python to replicate the results.

The tool can generate multiple years and create multiple numbers of claims in each year. We used 10 years with 50 claims each year in this study. This was chosen largely because of the size of data set and time limitations of processing in Excel, Python or R could support more simulations (Figure 1).

The simulation was not specifically based on an individual line of business or individual insurer's data. For simplicity, we used the Pareto distribution, which also gives us a good range. Claims were generated and triangles with annual development periods were created. A development pattern was applied. We were able to apply various types of inflation – economic, social and any one-off large spikes, to see how they would impact our assessment of the different methods.

Next, we wanted to start building up Scenarios increasing in complexity as we progressed (Figure 2).

Scenario A is very straightforward. It was one stable 4% rate of inflation for all the years we were investigating. We then moved to Scenario B where we included social inflation for the later years. Scenario C is still higher in complexity, with shock inflation – a spike of inflation within the year. Scenarios D1 and D2 went further, where we investigated how the frequency of claims could also impact our assessment of the underlying inflation. Scenarios A to D2 were all known when we

© The Institute and Faculty of Actuaries, 2025. Published by Cambridge University Press on behalf of The Institute and Faculty of Actuaries. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (https://creativecommons.org/licenses/by/4.0/), which permits unrestricted re-use, distribution and reproduction, provided the original article is properly cited.

https://doi.org/10.1017/S1357321725100196 Published online by Cambridge University Press

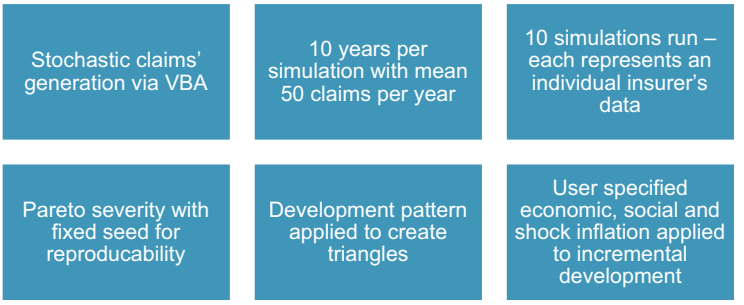


Figure 1. Estimation data.

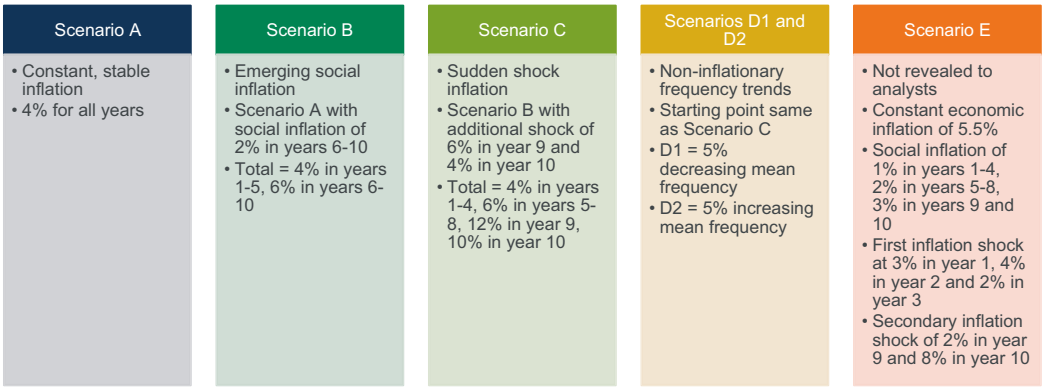


Figure 2. Scenarios.

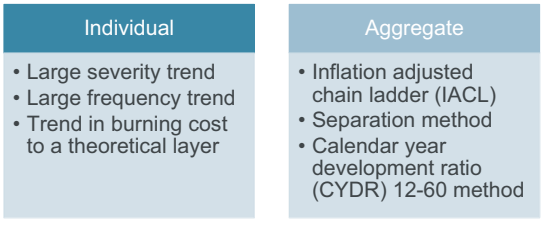


Figure 3. Individual versus aggregate claims methods.

did our assessment, that is, the information in terms of the input was known by those in the working group who were trying to test the various methods we had. Scenario E was slightly different. It was not revealed to the analysts what the inputs had been, to see how the various approaches would work on a blind test. It was also one of the more complicated scenarios, but it does give a feel for how well each of these methods may have worked in real life.

We had two groups of methods that we were looking into - individual and aggregate claims (Figure 3). By individual claims, we mean we were looking at how the severity trend was working for the large individual claims, the frequency trends based on the changes from inflation and the burning cost for a theoretical layer over each of these simulations. We were also able to condense those triangles into an aggregate triangle to have a look at the inflation adjusted chain ladder approach, the separation method and the development ratio on the calendar year development. I will pass to Erin (Bargate) now, who will discuss the results and the evaluation.

https://doi.org/10.1017/S1357321725100196 Published online by Cambridge University Press

Miss Bargate: We looked at the results of using the different methods on the data that we had generated.

For each scenario we have a chart titled "Output Comparison by Method- All Sim" that shows the average across all the simulations by year for each of the different methods compared to the input inflation. In this chart for all scenarios, the input inflation is the orange line and each of the other coloured lines represents the other different methods that were trialled. We are interested in identifying the line for one of the methods that matched the orange line representing input inflation as closely as possible. For each scenario, the chart titled "Inflation Estimation Results" was then generated using the method we thought performed best in the specific scenario. Each dot on these charts represents one of the ten individual simulations for that particular year. These can be thought of as an individual insurer's data at an aggregated level equivalent to the data of a reinsurer, consultant or broker.

In Scenario A (Figure 4), all the methods perform reasonably well. Granted, this was a very simple scenario where we just had constant 4% inflation, and all methods gave an output of 4% inflation. As we see, the coloured lines are all grouped around 4% inflation, except for the large frequency trend.

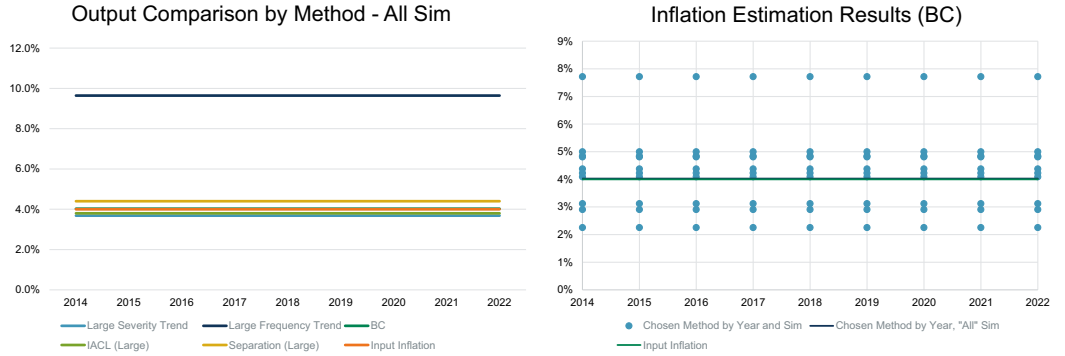


Figure 4. Scenario A. All methods perform reasonably well, although frequency trend an outlier in over-estimating.

Perhaps this is not too surprising. After some analysis on why the frequency method did not perform as well, we think it was purely an anomaly and can be attributed to the relatively small sample size, and also the judgement that is involved in selecting the threshold for the frequency. We did also find in general that, across scenarios, the frequency method tended to over-estimate the inflation, particularly when the inflation was low. Based on Scenario A, the burning cost method, which is the green line on the chart, most closely matched the input inflation.

For Scenario A, the chart "Inflation estimation results (BC)" shows each of the ten simulations for the burning cost method, with the green line representing input inflation and the dark blue the chosen method. The interesting thing to note is that the results show a wide spread, even though on average across the ten simulations the burning cost method performed quite well. One individual simulation selected a number nearly as high as 8% and another a number that was nearly as low as 2%. If we assume each of those dots represents an individual insurer, it shows that it is quite difficult for them to estimate what the true inflation might be if they are relying solely on their own data.

Next is Scenario B (Figure 5), which has constant inflation in the first five years and then an up-tick in six to ten years. On the chart "output comparison by method", the orange line is seen to be stable up until 2017 at 4%, then it jumps up to 6% from 2018 onwards.

https://doi.org/10.1017/S1357321725100196 Published online by Cambridge University Press

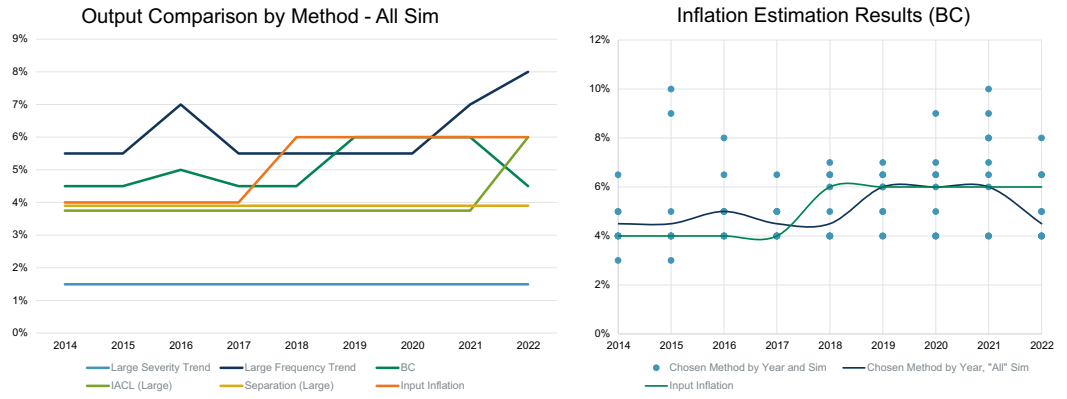


Figure 5. Scenario B.

In this case, we found that both the burning cost and the frequency trend (green line and dark blue line) were good at detecting the up-tick in inflation, whilst the other methods did not really detect that up-tick in the middle of the period. (See the yellow, the light green and the light blue lines that stay flat until 2021). Once again, the frequency method (dark blue line) slightly overestimated the up-tick, and so the burning cost method (darker green line) was once again the best performing in Scenario B. The chart "Inflation estimation results (BC)" again shows a significant amount of variability in individual results. This highlights how important it is for the judgement that needs to be used while making selections and the threshold levels required to make these selections, particularly while working with a small amount of data as an individual insurer might be.

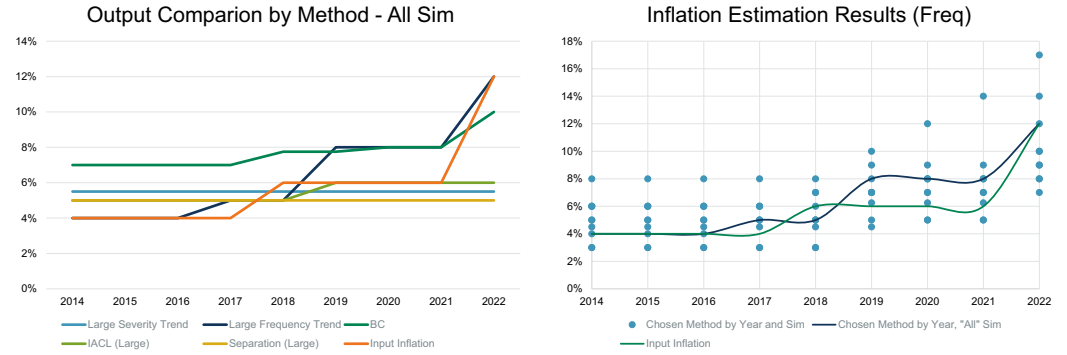


Figure 6. Scenario C. Frequency trend found to be most responsive with regards changing inflation in period.

Next in Scenario C (Figure 6), we have Scenario B plus a further up-tick in inflation towards the end of the period: 6% to 12%.

In this case, we saw that the burning cost and frequency methods performed relatively well. However, we thought the frequency method was better because the burning cost method overstated the inflation at the beginning of the period, whereas the frequency method was closer to the input (orange line) at the beginning and then accounted for the significant up-tick at the end.

What is interesting in this case is there is not as much of a spread in the individual simulations for the frequency method. Another interesting observation is these methods are similar, yet here we had the severity method (for example, the blue line and some of the other lines) perform better than they did in Scenario B. We concluded this was unlikely to be because of simulation error,

https://doi.org/10.1017/S1357321725100196 Published online by Cambridge University Press

because we used a fixed seed and it was more likely to be because of users' judgements. However, the highest selection would have still been close (16% to 17%) and would still have been a significant overstatement in the 2022 period; while the lowest would still be around 7%, a significant understatement. What is also interesting to note is that although there was still a range towards the end of the period, every simulation did pick up that there was an up-tick in the trend and that was not missed when using the frequency method.

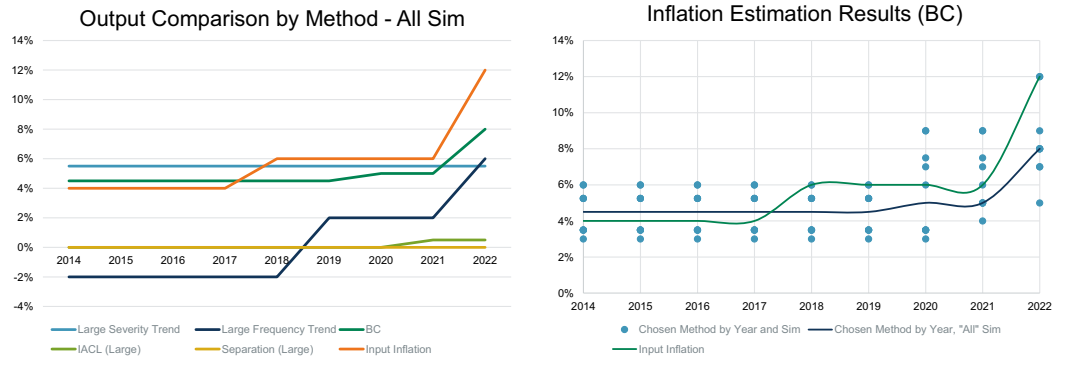


Figure 7. Scenario D1. Aggregate and frequency methods were poor at distinguishing frequency trend from claim cost (severity) inflation.

In Scenario D1 (Figure 7), we applied a negative frequency trend in addition to what was used in Scenario C. In Scenarios D1 and D2, we try to assess how well the various methods can essentially look through the frequency effect to ascertain the underlying claims cost trend. The input inflation in Scenario D1 (the orange line) is identical to Scenario C, because we use the same claims cost trend, to be able to look through frequency trend to identify what the true severity trend. In this case (Scenario D1), the methods were not very good at dealing with the frequency trend. We can see that the aggregate methods and the frequency methods were particularly poor at matching input inflation. Once again, the burning cost method was the best performer, although the severity method (blue line) matched quite closely on average, the burning cost method (BC, green line) allowed us to pick up the uptick in the end.

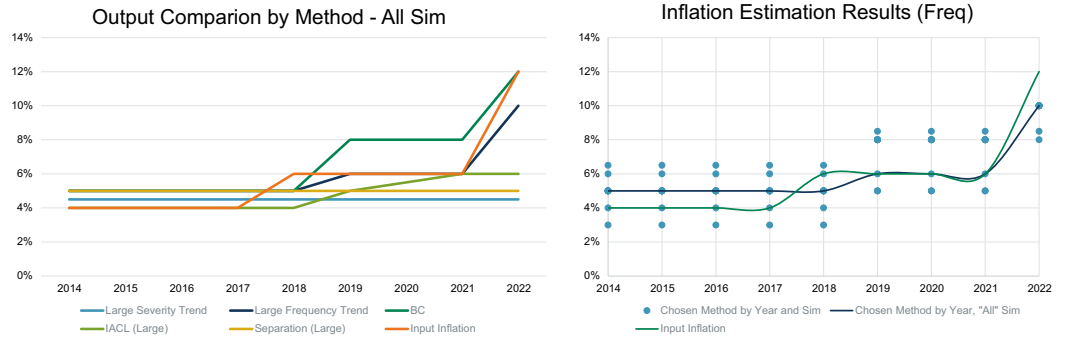


Figure 8. Scenario D2. Reasonably consistent performance across all methods and more robust at dealing with rising versus falling frequency. Frequency trend method best performer.

https://doi.org/10.1017/S1357321725100196 Published online by Cambridge University Press

We move on to Scenario D2, where we have a positive frequency trend instead of a negative frequency trend. In general, we found that the methods seemed to handle this much better than the negative frequency trend. You can see that most of the other coloured lines are much closer to the orange line (input inflation). What is interesting is that while we expected the frequency method with the positive frequency trend might overstate the claims inflation and have an amplifying effect, this does not seem to have been the case. This method was the best performer in terms of most closely matching the true input inflation. There is some variability in the inflation estimation results for the frequency trend (see right-hand side of Figure 8), but the spread is probably a little bit tighter than in some of the previous scenarios.

In all the Scenarios A to D, the analyst knew what the input inflation had been, which may have slightly skewed their judgement or biased them towards certain selections. That would not be the case in the real world. What we wanted was to try to test what we thought was a more real-world scenario, where the user who was trying to select the inflation would not have that prior knowledge, and see which method in that case might help them to determine the true, underlying inflation.

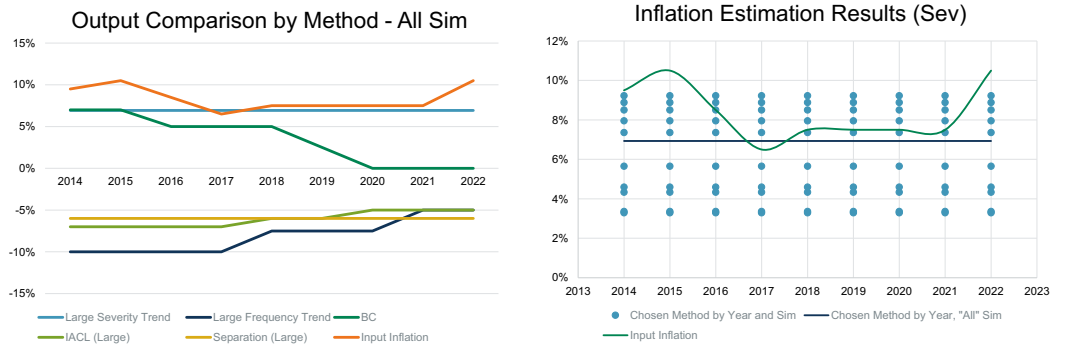


Figure 9. Scenario E. Broad range to be expected, given varying, unknown parameters. BC overcompensates for decreasing inflation. Severity trend, perhaps unsurprisingly, most adept at picking out long-run claim cost trend.

It is worth noting that Scenario E (Figure 9) did have a negative frequency trend. As we saw with the results in Scenario D1, none of the methods were very good at picking up the true underlying severity inflation trend with a negative frequency trend, which may also have influenced why some of them did not perform very well. The aggregate methods (light green and yellow lines) did not perform well, perhaps because these aggregate methods tended to conflate the frequency and lost cost trends. In terms of the methods that did come close to matching the orange line of input inflation, we have the burning cost method and the large severity trend method. The burning cost (darker green line) did quite well early on, but it underestimated the inflation in the later years. Once again, it was over-compensating for that negative frequency trend and not picking up the severity trend. In this case, we have the large severity trend performing best (straight blue line). In this scenario, that was best at picking up the long run trend, which is perhaps at odds with some of the earlier scenarios. However, in the chart on inflation estimation results, there was an extremely wide range in the individual simulations. It is also worth noting that when we were completing this exercise and applying that large severity trend, we selected a single number for the severity to be used as an initial gauge, which is why in the figures it is just a

https://doi.org/10.1017/S1357321725100196 Published online by Cambridge University Press

straight blue line. However, if you were implementing this in practice, it could also vary by year in the same way as the frequency and burning cost methods. Potentially that might have meant, in some of the other scenarios, more variability in the inflation and could have matched more closely.

Now that we have seen the results of the exercise we performed, I will hand back to Cian (Creedon).

Mr Creedon: So far, we have talked about a cohort of historical claims data, what methods and approaches are most suitable to pick out the claims trend, and the claims inflation in that cohort of data. This has been done under a variety of different scenarios designed to represent the increasing complexity of the real world, and for a variety of different approaches, which Erin (Bargate) has evaluated.

However, even with our best attempts to replicate the real world with pseudo-data, the real world is a more complicated place. There are various ancillary considerations and practical reasons for these considerations. In terms of ancillary considerations and claims inflation estimation, inflation gearing is a reasonably well-known phenomenon. That essentially means that, if we have a gross level of claims inflation, say X%, the inflation of claims excess a certain threshold, say into an insurance layer or into excess insurance layer versus the primary layer, will typically be greater than the underlying gross, or ground-up inflation. If the ground-up claim severity follows a Pareto distribution, then the level of gearing, or the geared inflation excess the threshold, or the geared inflation into the re-insurance layer, is independent of that threshold. This can be calculated by a simple formula. The geared inflation is equal to one plus the original inflation to the power of the Pareto alpha, that is:

$$i_{geared} = (1 + i_{FGU})^{\propto}$$

Where $\propto$ is the shape parameter of the fully ground up (FGU) Pareto severity distribution.

However, that is only a special case where the severity follows a simple Pareto distribution. In the case that it follows other distributions, then the level of gearing excess the threshold does depend on the threshold selected.

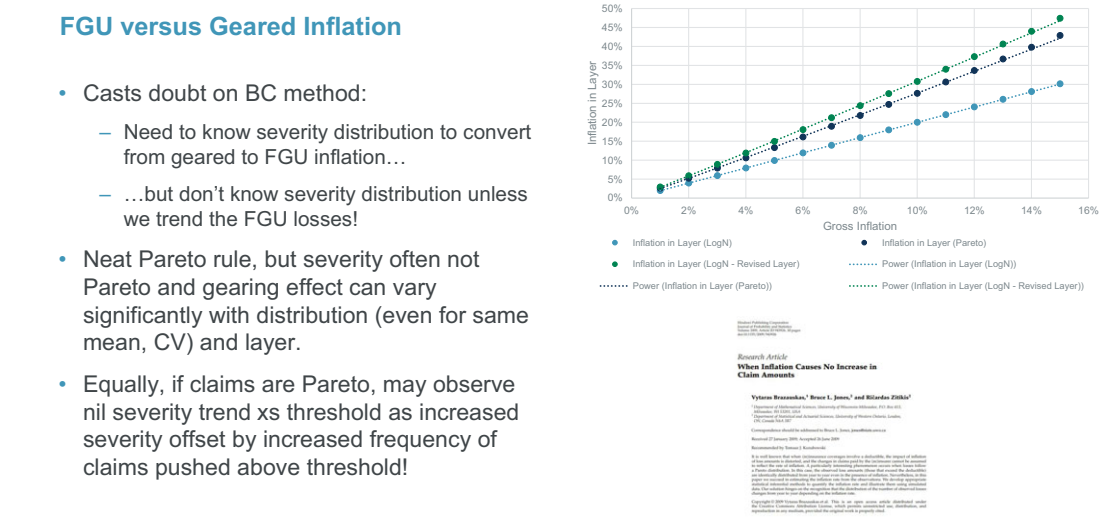


Figure 10. Fully geared up versus geared inflation.

Figure 10 shows the level of geared inflation observed for a normal distribution under two excess thresholds and a Pareto distribution, and how they differ from each other. In particular, the lognormal and Pareto distributions are seen to have the same mean and standard deviation. They

https://doi.org/10.1017/S1357321725100196 Published online by Cambridge University Press

are equivalent in terms of the first two moments, but obviously the lognormal will have a different shape than the Pareto. There is quite a significant gap between the lines, represented by the lognormal distributions with different thresholds (green and the light blue lines). In other words, 8% gross inflation can lead to a 15% geared inflation excess one threshold, or 25% geared inflation excess another threshold. Thresholds were chosen to be one in 2 million units, so a factor of two apart. Another corollary to this point is that, when inflation causes no increase in claims amounts, as in the case that claims do follow a Pareto distribution, we may observe nil severity inflation excess the threshold for an excess or reinsurance layer vis a vis FGU.

The reason is that as gross claims are inflated an additional number of claims are, essentially, pushed above the threshold, just as claims that are already excess threshold for analysis are inflated further. In other words, the inflation due to claims already in the layer excess of threshold increasing is offset by the increased frequency of claims that were previously below the threshold. If data is left truncated, then there may be nil severity inflation observed, even though there will be gross inflation and inflation of aggregate claims amounts. The aggregate claims will be inflated even though there is no severity trend observable in excess of the threshold.

A nuance that has become increasingly well understood over time, and certainly since the inflation spike of 2021, is the distinction between settlement and origin year inflation. All the methods we have discussed so far have been conducted on claims triangles or individual claims that have been generated according to an origin year basis. In general, when we do actuarial work, be it reserving, pricing, capital modelling or other analyses, we tend to think of claims in origin year terms. Origin year could be accident year or underwriting year. In other words, we are concerned with the reserves associated with a certain underwriting year, the rate change associated with that underwriting year and the prior loss ratio certification for that underwriting year.

However, when we think about claims inflation or economic inflation that tends to be measured on a calendar year basis. The change in CPI associated with any one year is the change in the price of a basket of goods from one calendar year to another. The claims from a given origin year (accident year or underwriting year) will typically be paid over several different calendar years. The accident year inflation, or the underwriting year inflation, can be thought of as a pattern-weighted blend of the observed calendar year inflation over a number of years. Depending on how long or how short the payment pattern is, or the settlement pattern, that increased length then increases how much further to predict into the future.

One thing to note on this is that claims can effectively be settled where the payment is agreed more quickly than they are paid. For example, a reinsurance claim is not paid until the underlying insurance claim is paid, and the underlying insurance claim may be settled some time before it is paid. So, using a payment pattern to calculate the blend of future calendar year inflation may be looking slightly further into the future instead of weighting slightly more heavily towards earlier origin years.

Consider an example. If we are in 2023, let's say, and are calculating the claims inflation associated with 2023 origin year and 10 previous origin years, we might chose the payment pattern shown in Figure 11.

The inflation associated with an origin year as at 2023 is attributable to historical versus future calendar years, that is, while reserving or pricing in 2023, inflation for the 2023 underwriting year is, essentially, a blend of the next ten calendar years of inflation (Figure 12).

We have talked at reasonable length about how to estimate inflation present in claims data, be it one's own data, external benchmark data, or be it a data collated by a reinsurer or broker from several companies. Often when claims inflation is considered, we turn to an external index to map the inflation, or to give an estimate of claims inflation which to apply. An external index can assist in validating empirical estimates or replace them when an empirical estimate is unfeasible. The advantage of taking a claims inflation estimate from an external index is that it is independent and does not necessarily rely on judgement, so therefore it is perhaps less open to challenge, and it is easy to communicate. In the absence of access to a lot of external claims data, this is much easier to

https://doi.org/10.1017/S1357321725100196 Published online by Cambridge University Press

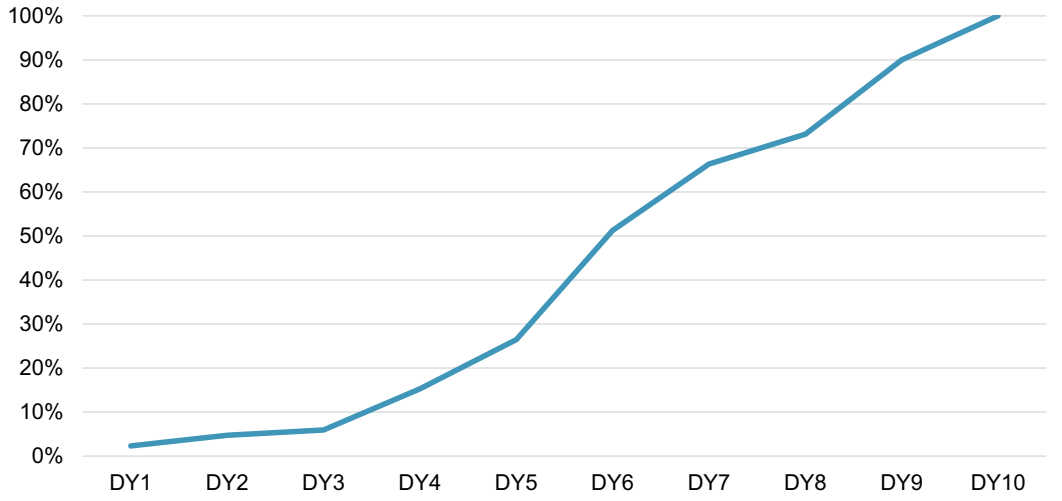


Figure 11. Example cumulative payment pattern.

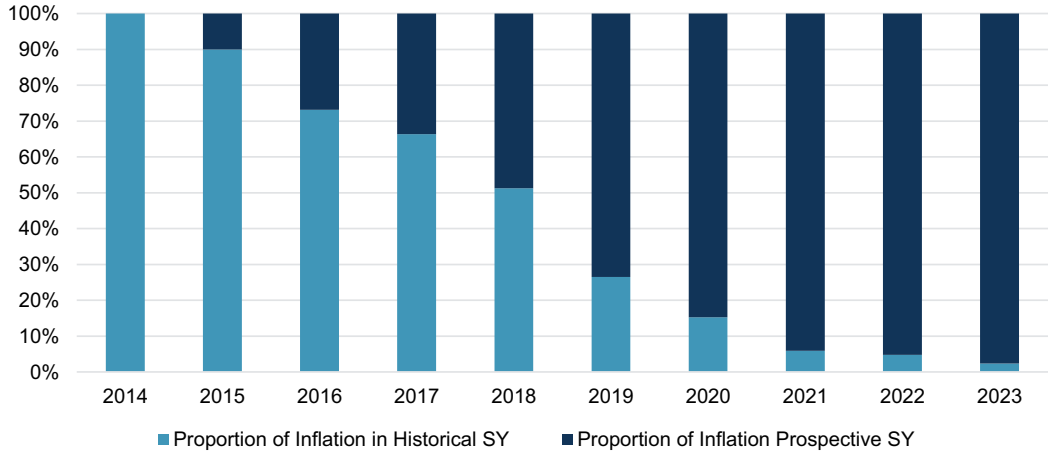


Figure 12. Origin year inflation split – historical settlement year versus prospective settlement years.

get and much more usable. The inflation pick need not be tied to a single index. More than one can be used and blended. Adjustments in the form of loadings can be applied. The independence point is quite helpful: in situations where there is no access to data such as a reinsurance broker, using an external index can be useful. For example, to explain that hull inflation is increasing, it is reasonable to use steel price index movements since hull manufacture and repair costs would be tied to this very closely.

One the disadvantages of using external indices is that it is not always evident what external indices are most relevant to the class and business in consideration. Consider the reasonably simple case that you are analysing motor accidental damage. Motor accident claims costs can be tied to used-car price indices. As you move into more complicated lines of business, such as residential property, household claims inflation can typically be tied to a rebuilding cost index. But a rebuilding cost index might not be available, so a blend between labour and materials can be used. But questions arise as to which indices to blend –- what kind of labour, what materials and so on. As well as this, in using a calendar year economic index we must map that back to underwriting year, and it can lead to a false sense of security. For instance, in a situation where an

https://doi.org/10.1017/S1357321725100196 Published online by Cambridge University Press

external index is not rising, it may not necessarily mean claims inflation is not rising without considering factors driving claims inflation.

If a model can explain the inflation present in claims data via some external indices, those external indices can be updated quite rapidly, or will become available reasonably rapidly, and so it becomes possible for the origin year inflation indices to be updated more rapidly to reflect changing external inflation.

As mentioned, choosing the inflation indices to explain claims inflation can be challenging, but there are a couple of approaches: a mathematical approach or a first principles logic-based approach (Figure 13). An example of using the first principles approach could be that household claims inflation could be a blend of labour and materials. Here is an example of a logic-based approach from a study for AON to estimate the inflation associated with cargo claims. The logic used in estimation is as follows:

- We first ascertain the proportion of claims that occur on land, in situ and in transit and at sea. In this example, just under 50% of cargo claims occur when the cargo is in situ and a further 40% occur when the cargo is at sea and the remaining when on the road and elsewhere.
- Next, we try to ascertain the types of goods – finished goods, crude oil, agricultural products and so on.
- We obtain relevant inflation indices for each of those products, and we use the breakdown of claims plus what is being carried to come up with an appropriate blend.

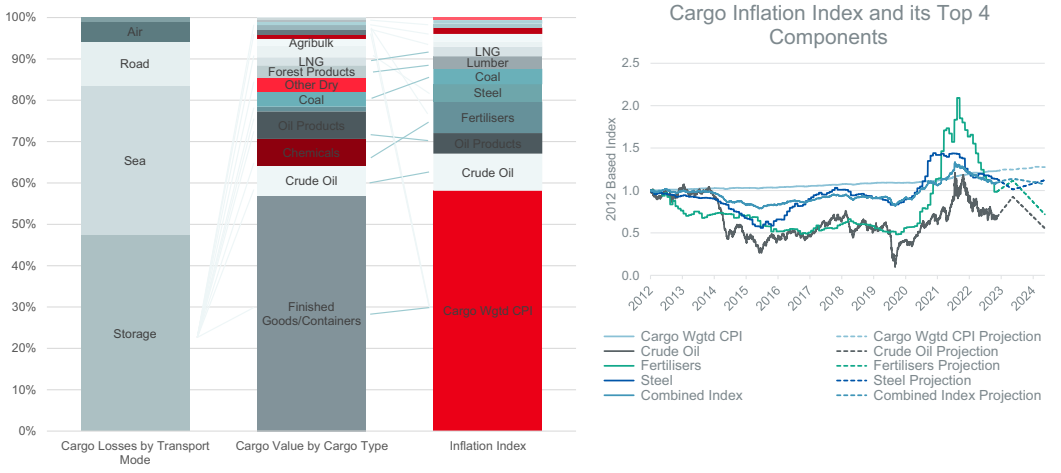
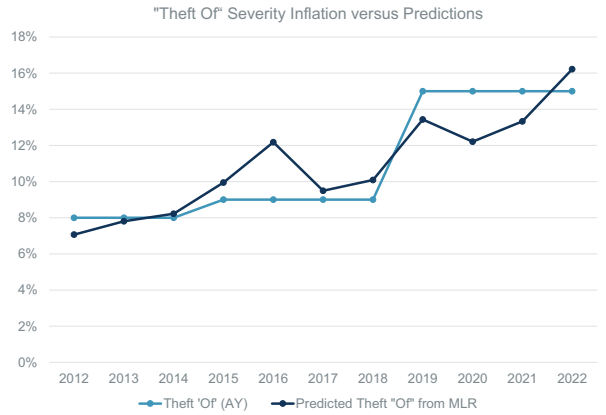


Figure 13. First principles approach.

This kind of logic-based approach can be applied to any class of business. For instance, for inflation associated with renewables, the renewables claims will either be for repair, total loss or business interruption. We would need the proportion of claims associated with each type of damage, and the drivers of claims associated with each. This approach is one that a non-actuary can challenge as well, and so it is quite useful. However, one could also use a more mathematical approach, namely multilinear regression, to map measured claims inflation to a variety of different indices.

Here is a test case we applied whereby we estimated settlement year inflation associated with vehicle theft severity in the UK (Figure 14).

https://doi.org/10.1017/S1357321725100196 Published online by Cambridge University Press



- Estimated vehicle theft claim severity inflation from data provided (kindly) by ABI
- Obtained UK CPI and PPI breakdown and first assessed correlation of components against fitted claims' inflation.
- Metal production inputs, earnings and tobacco (!) found to be heavily correlated.
- Accordingly, applied multi-linear regression model of these indices against estimated inflation.
- Surprisingly successful fit!

Figure 14. Fitting approach – initial experiments.

The data is from the Association of British Insurers. We built a model that explained or matched severity inflation via three economic indices relevant to the UK, namely metal production (cars are made in the large part of metal); earnings (repair costs to fix your cars); and, more bizarrely, tobacco price. Those three indices, including tobacco price, can be used to provide a reasonably good fit to motor vehicle theft severity in the UK. This approach could be applied to any line of business, using various available indices and a multilinear regression model.

We looked at why we need to estimate the inflation in claims. Prior to the COVID pandemic (2020–2021), economic inflation in western economies was generally low and stable. However, as we came out of the pandemic, supply chain shortages led to a spike – a period of supernormal inflation that had not been seen since the '70s and '80s.

This prompted a scramble to ensure pricing and reserving accurately reflected the changing economic environment. However, before we could undertake those exercises, we first had to assess inflation present in their historical claims and how that was impacted by economic inflation, as a starting point to begin to correct pricing, reserving and risk adjusted rate change. Essentially, what our working party considered is how to amend or conduct the reserving exercise the light of inflation volatility, or when there is a sudden change in inflation, to ensure that reserve redundancy is as low as possible, or reserves most accurately reflect the future payments that are to be made (Figure 15). There are four approaches to be considered:

- Loading priors used in reserving.
- Inflation adjusted chain ladder.
- Uplifting the cash flows that are implied by the reserves set
- Applying explicit management loadings.

We are not the first Claims Inflation Working Party. There was a prior Claims Inflation Working Party, which was run by Simon Sheaf, Simon Brickman and Will Forster, about 20 years ago. One of the first things they remark in their paper, which deals heavily with inflation estimation, is that it is hard to do, notoriously so. We have come to echo those comments. There is no clear automatic method to estimating inflation in historical data. It requires quite a considered degree of actuarial judgement, as should any actuarial exercise. It requires quite a large amount of work, and the data required is quite significant. A single insurer may not necessarily have sufficient data to be able to separate general noise/ randomness from trends in claims inflation. Although we found the burning cost method to be quite a useful or powerful inflation estimation method, the impact of inflation gearing in excess of the threshold throws it into question. There are various other challenges as well, as have been alluded to in the paper.

https://doi.org/10.1017/S1357321725100196 Published online by Cambridge University Press

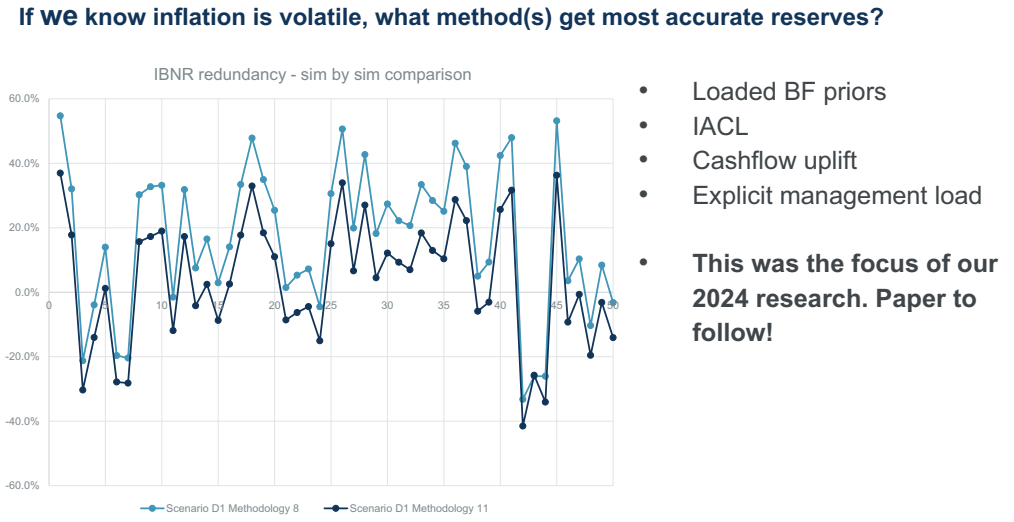


Figure 15. Reserving.

In addition to our firms, who kindly give us the time to pursue this hobby, there are quite a few people who have contributed to the work we have done in one form or another and we are very grateful to them.

Mr Lenney (starting the Q&A session): Is there a reason you did not allow the large severity trend method to vary by year?

Mr Creedon: No, there is not a reason. We were considering about six methods in the work. That was six methods across six or seven Scenarios, so it was times ten simulations average. There were a lot of iterations of judgement to be done. When we initially set up the work, we naively did not think that the large severity method would be as powerful as the other ones and simply did not set things up to have it varying by year. That was an oversight. There is no reason why, just like the frequency method, you could not set it up as an index and have large severity trend varying by year. In fact, you almost certainly should.

Miss Bargate: I have found that sometimes while working year by year on an individual basis, there is substantial volatility in the estimates, particularly if it is a small data set and there are lots of different claims. We usually select a long-run trend instead of trying to precisely match it to year-by-year; particularly if, for example, it is calendar year inflation or settlement year inflation that we are working on, and then we apply it to origin year. That is going to blur the line even further and introduces more volatility. Sometimes looking at the long-run trend is helpful also, in addition to considering the other methods.

Question: In the first part of the presentation, you shared that certain methodologies were best, albeit highly variable for each type of Scenario that you created. Did the method that you decided was best having performed the analysis match the methodology that you thought would be best from first principles, given the Scenario that had been created? For example, did burning cost outperform even when it should not have? If so, why do you think that was? If another method should have performed better, why do you think it did not?

Mr Lenney: There were some obvious outcomes. The aggregate methods did not perform as well for the more complicated Scenarios that we put forward. I guess that was something we were expecting to see. Perhaps, the burning cost did better than I thought it was going to do, in terms of some of the outcomes that came through.

https://doi.org/10.1017/S1357321725100196 Published online by Cambridge University Press

Mr Creedon: Naively, we expected that because the aggregate methods use aggregate data, you would not expect them to be as accurate. If frequency is changing, aggregate and frequency methods will blur frequency and severity effects, which is something of a worry. One important thing to consider is if frequency is reducing but claims are more severe – are you in a situation where claims are getting more expensive when they do happen, but they are happening less often? Those are very much real-world situations. You may be working on motor accidental damage, for instance. As more sensors and cameras and digital gadgets get added to cars, people tend to crash them less frequently. They certainly tend to have fewer minor parking claims. But when claims do happen, they are more expensive. A plastic fender is reasonably cheap to fix. A fender that is full of cameras and sensors is a lot more expensive. The question to ask is if there is a reason that frequency and severity might be moving in opposite directions, or at least exhibiting separate trends. Did methods perform better than we expected to and why? In the Scenario where we had rising claims frequency, methods were better at picking that up or better at allowing for that than they were a falling claims frequency. It is good to apply a number of different methods if possible.

Question: How do you bridge the gap of settlement year versus underwriting year or accident year, as reserving and planning are mostly done on an accident year basis, whereas indices are more aligned to settlement years?

Mr Creedon: In that case, we use a payment pattern to create a blend. If index inflation is on a calendar year basis, use a payment pattern to blend those calendar year inflation estimates onto a large year basis. The only problem is the need to determine an estimate of future calendar year inflation, which is fundamentally a guess, if you are dealing with the very long-term tail line of business, maybe even guessing twenty years into the future. It will be going back to relying on long-term averages, with a bit of judgement involved. It is a reasonably quick and easy thing to do once set up.

Question: How did you go about selecting 5 million excess 2 million as the reference layer for the burning cost method?

Mr Creedon: It was purely based on the mean of the large severity distribution. It was set somewhere above the mean. The result can vary based on the layer of threshold selected. It is something to bear in mind and it is worth trying different thresholds, different pseudo layers and so on until you find the combination that leads to reasonably stable estimates.

Question: For each of the Scenarios, was the cumulative difference between the input inflation and the chosen method calculated?

Miss Bargate: I don't think we did this. The chosen method was picked by looking at the graphs and making a decision, rather than looking at the cumulative difference. This was already quite a labour-intensive exercise to get all the different estimates on a different basis. We thought making the selections based on the graph was the most straightforward and pragmatic way to do it.

Mr Creedon: If you are in a work Scenario and not just doing this for research purposes, then given how important some of these numbers could be in setting reserves, it is the sort of thing to get committee sign off on, or at least consensus on. As the question alluded to, you may want to present metrics of goodness of and fit about cumulative difference before you come to your chosen estimate, if not chosen method.

Moderator: Thank you again to the IFoA for facilitating the dissemination of the paper, for considering us as runners up for the Brian Hey prize and for assisting in running this webinar.

Cite this article: Practical considerations in claims inflation estimation. *British Actuarial Journal*. https://doi.org/10.1017/S1357321725100196

https://doi.org/10.1017/S1357321725100196 Published online by Cambridge University Press