

SESSIONAL PAPER

Claims inflation estimation: a practical guide for historical data

Cian Creedon, Erin Bargate, Shane Lenney, Marcus Schofield and Richard Stock

Institute and Faculty of Actuaries, London, UK
Email: ocriodac@tcd.ie

1. Impetus and Motivation

Claims inflation is one of the key assumptions used by non-life actuaries. It is relevant to virtually every aspect of their work, whether reserving, pricing, planning or capital modelling. In some instances, an appreciation of past inflation is needed, whereas in other cases, it is the estimation of future inflation rates that is key. Unfortunately, claims inflation is notoriously difficult to measure with any degree of certainty (Sheaf et al., 2005)

Over the past three decades, central banks in Western economies have sought to maintain a low and stable level of inflation, as measured by relevant price indices. By and large, this aspiration was successful until the year following the outbreak of the COVID-19 pandemic in 2020. From roughly the second half of 2021, price inflation in Europe, the US, the UK and most other nations exhibited a sharp increase. Although price inflation reached its peak in 2022 (and has since reduced significantly), it remains highly elevated in many economies, vis à vis their central bank targets; as of the time of writing this paper (Q1 2024).

With certain exceptions (e.g. general or punitive damages), the aim of a general insurance claim is to rectify a financial loss. As such, general economic or price inflation can be expected to have a significant impact on the level of claim pay outs. This impact, however, can vary markedly between line of business and territory and may temporally lag or, indeed, occur in advance of periods of significant change in general economic inflation.

Given the link between price inflation and claim payments, periods of supernormal economic inflation have the potential to impact the solvency of general insurance firms through under-provisioning, under-pricing or under-capitalising. The Prudential Regulation Authority (PRA) of the UK thus listed inflation's effect on financial resilience as a key focus of general insurance supervision in their letter to CEOs in January 2023 (PRA, 2023). It is worth noting also, that claims inflation remains a supervisory priority for the PRA as of January 2024 (Sheaf et al., 2005)

Prompted by these concerns, a group of like-minded actuaries came together at the beginning of 2023 to discuss approaches to managing the impact of inflation uncertainty, as it pertains to general insurance. The resulting discussions indicated a lack of definitive and easily comprehensible guidance as to this problem. As such, the individuals formed an actuarial working party on the subject.

The aim of the Claims Inflation Working Party is to produce a variety of practical and digestible guidance papers, each dealing with an individual aspect as to how to manage inflation uncertainty.

2. Abstract

The goal of the working party is to produce practical and digestible guidance. Accordingly, we endeavour to be as clear and concise as possible when discussing techniques and concepts, with as many examples provided as feasible.

In Section 4 we outline our understanding of the term “claims inflation”. We will also allude to some of the pitfalls that may be encountered in endeavouring to measure it.

In Section 5 the data used in all the estimation research we conducted; as well as the assumptions and approach used to generate that data. Full information on the data generation approach is deferred to Appendix A. We also discuss how we used these data to create a variety of “scenarios” under which to estimate inflation and how the data were considered conceptually in the estimation work.

Following on from this point, we will cover, in very general terms, the studied set of claims inflation estimation methods we have employed in our work. These can be categorised into methods applicable to aggregate claim sets (i.e. run-off triangles) and methods applicable to individual claims. Additional information on these methods may be found in Appendix B and papers referenced.

Having provided an overview of the techniques and the data sets on which they were used, we next evaluate the inflation estimation methods used by each scenario. This evaluation includes a comparison of “input” against estimated inflation, as well as practical challenges and insights.

Section 11 of the paper considers some of the ancillary issues that arose during the estimation work, their practical implications, and ways to manage them.

Finally, we consider the next steps and areas of research for the Working Party.

3. Introduction

In June of 2023, the PRA issued UK general insurance chief actuaries with a follow-up report to their October 2022 thematic review (Bohnert, 2015). In this report, it was noted that many insurers had yet to observe recent supernormal economic inflation manifest itself as claims inflation, particularly in third-party casualty lines. Although claims inflation may appear to be disconnected from economic inflation in many lines of business, this is considered to be unlikely (given insurance’s purpose in rectifying a financial loss). Rather, it is likely that claims inflation may lag economic inflation for certain lines.

Although a total disconnect is highly unlikely, the link between economic and claims inflation can vary from being reasonably apparent (such as in first-party motor property) to less apparent (such as in medical malpractice). Accordingly, the question of how (and indeed, when) best to allow for the effects of supernormal economic inflation (and periods of heightened claims inflation generally) in general insurance actuarial practices becomes challenging.

Indeed, in order to answer the question as to how economic inflation impacts claims inflation, we first need to achieve a solid understanding of what claims inflation has been present in our data historically. This, then, becomes the pivotal focus of this paper, namely:

For a given cohort of claims data, how do we best gauge the level of historical claims inflation?

In line with the stated purpose of the Working Party, the objective of this paper is to seek to provide an introductory guide to actuarial practitioners in their attempts to answer this question. In this paper, we aim to evaluate and provide guidance on some common methods which may be used to estimate inflation present in historical data.

4. Definition of Claims Inflation

4.1 General

Before proceeding further, it is worth noting that the term “claims inflation” may be understood quite differently by different practitioners in general insurance. As such, it is considered worthwhile to define what is meant by claims inflation for the purposes of this discussion. Given that much of the Working Party’s experience is in the London Market, we have chosen to adopt the definition provided by Lloyd’s of London (Lloyd’s of London, 2022a, 2022b), namely to define claims inflation:

“... as the change in claims cost of a like for like policy over time.”

Here, the term “claims costs” incorporates all costs associated with settling a (re)insurance claim, including claims handling/allocated loss adjustment expenses.

Claims inflation can be further broken down into the sum of its economic and excess components. The economic component can be thought of as the element of inflation directly attributable to one or more economic factors, which may be captured by relevant, published economic indices (noting there is no prescribed “correct” index for any line of business or territory). Excess inflation is then the difference between this published index component and total inflation.

“Social” inflation is by no means a new term, but has seen increased prevalence of use over the last decade or so; particularly with regard to Casualty lines. It is encompassed within the excess inflation element and, more specifically, per Lloyd’s of London:

“... narrowly pertains to claims inflation as a result of societal trends.”

It should be noted that the overall definition of claims inflation provided above encompasses both frequency and severity effects; i.e. an increased incidence of claims occurring can be considered claims inflation (in that claims in the aggregate will increase); as will an increase in claims costs with no change in incidence rate.

Equally, it is worth recalling (Sheaf et al., 2005) that increased claims costs (i.e. severity) can manifest themselves as a trend in frequency if one is considering non-primary or non-proportional reinsurance layers. Furthermore, although both trend in frequency or severity can be considered to be inflation under this definition, separate consideration of these two components is often crucial; as it will inform considerations such as policy structure (e.g. deductibles), reinsurance purchase, etc.

4.2 Specific Scope of this Paper

As is discussed in subsequent sections, the worked examples and scenarios discussed in this paper have been constructed using artificially generated pseudo-data. As such, consideration of inflation estimates will encompass both excess and economic components combined. Equally, as the impetus for the formation of the Working Party was a spike in general, global price inflation, much of the focus of this paper relates to severity trend inflation.

4.3 Words of Caution

Over any concerted period of time, much can change for an insurance company. Even within a class of business, mix may change as new territories and sub-lines are entered, perils emerge, and staff change. Terms and conditions may change, such as exclusions or limits/deductibles. Reserving philosophy/strength, too, is liable to some movement over time, perhaps simply due to

more efficient reporting. All these effects may hinder any efforts to estimate historical claims inflation, even if perfectly well known to the analyst performing the estimation. Even should the historical business be relatively stable and homogeneous, challenges may arise from the reserving philosophy, such as very low opening reserves or maintaining low reserves until time of settlement, for example.

Take, for instance, the simple case of an increase to original policy limits. It is quite likely that we will see an increase in observed average severity following this change, which we may mistake for an increase in “true” severity (i.e. an increase in severity of uncapped claims). However, even fully armed with the knowledge that this change took place, allowing for it will be challenging. We likely cannot, for instance, as-if historical claims to a higher limit (the data is effectively censored/truncated). Equally, capping of newer claims at the old limit will be somewhat questionable, as the increased limit may have led to a change in exposure mix (higher limit may appeal to different insureds).

None of the above is to admit defeat in the task of historical claims inflation estimation. Rather, we caution the user to avoid doing so blindly. In providing worked examples based upon pseudo-data, we have endeavoured to strike a balance between realism and achieving sufficient simplicity such that the workings of the various methods used are readily apparent.

As actuaries, particularly those of us who work in reinsurance or the London Market, inconsistent, inconstant and, at times, poor data is something of a fact of life. In any actuarial exercise, a degree of data judgement and adjustment is likely to be necessary and inflation estimation is no different. A rather pithy and reasonably relevant discussion of some such concerns may be found in Flower *et al.* (2006)

5. Data Used in Estimation Work

5.1 Overview

For the purpose of investigating the application of inflation estimation methods, it was desirable to create a data set that was at once easy to work with; easy to reproduce; and reasonably representative of “real world” data. Accordingly, a tool was constructed in Excel/VBA to generate claims for a pseudo class of business, mirroring capital modelling underwriting risk thinking; i.e. aggregate attritional claims and individual non-aggregating large claims.

These claims would be generated over a number of simulated years in each given trial, with inflation and development applied to construct pseudo-real triangulations of individual claims; which would then be combined into aggregate triangles. The parameters used in generating the data were originally based on real-world data, but with anonymity of origin retained throughout and amongst the working party.

The impetus for splitting the data into attritional and large claims is that, in addition to its being a common conceptual framework, as mentioned above, inflation drivers often align closely with these groupings. For instance, in the case of a comprehensive motor policy, smaller attritional claims (typically first-party vehicle damage) may be simply and directly linked to vehicle parts and mechanic labour; whereas drivers of trends in large (typically bodily injury) claims will be substantially more complex and nuanced (e.g. closely linked with medical costs).

In addition, a number of more in-depth claims simulation tools and frameworks exist in the public sphere that could be used for the purposes of generating synthetic claims data. Examples of these include “SynthETIC” (Avanzi, 2001) and Gabrielli and Wuethrich’s “Claims History Simulation Machine” (Gabriella & Wuethrich, 2018). However, the authors have chosen Excel/VBA as a simulation tool owing to its ubiquity and relative ease of use. The tool used by the authors will be made readily available to readers and, it is believed, may be employed and examined with minimal additional specialist knowledge. We would, though, encourage the interested reader to explore the tools listed above, should they desire.

The workbook used to generate the data may be provided upon request.

5.2 Conceptualisation

In conducting the estimation work, the Working Party found it helpful to rationalise, at a very high-level, the theoretical provenance of the data used and who might be holding it.

In particular, (as we shall see below), the modest development tail and comparatively high level of inflation reminded some of us of claims made, professional indemnity business. Equally, with a mean frequency of 50 large claims per annum per simulation and ten such simulations, one could further conceptualise each individual simulation as a real-life large claims data set from a single insurer; with the holder of all ten data sets potentially being a reinsurer, reinsurance broker or consultancy.

In reality, any given insurer will have comparatively few claims of their “own” with which to estimate inflation and so the concept of a single simulation representing the full, large data set available to a single insurer for a given class of business makes intuitive sense. Separating out true inflationary effects from random variation is a very real challenge.

This potential lack of in-house data and the inherent challenge it presents vis à vis signal versus noise may prompt individual insurers to seek external data to aid in estimation efforts. Such data sets may be available from business partners (such as reinsurance brokers or consultants) or may be purchased from an external provider.

The London Market Association produces (for the use of its members) claims and premium triangles across Lloyd’s members, at a reasonably granular line of business (risk code) level, which may be useful for the application of aggregate estimation methods. At a more specialist level, providers such as Zywave, Stanford Securities Analytics or others besides may provide useful, individual claim data. As with any external data source, careful consideration should be given as to its applicability.

5.3 Generation Method

VBA was used within Excel to generate:

- Ten simulated years of individual large claims, aggregate attritional claims sets and event type claims per trial
- A random number of claims in each accident (equivalently underwriting) year
- Each claim having random size.

These claims were initially generated on an ultimate basis without inflation. Inflation was then applied on an incremental development basis to obtain a set of ultimate, inflated claims and claim triangle.

Functionality was provided in the tool to allow for three different inflation indices, which we dubbed economic, excess and social. The inflation parameters specified in each index were then summed in any given year to produce a total inflation value. Although this summation step is effectively equivalent to using a single index, we found the split to aid in conceptualisation and in discussing the scenarios of interest.

A full description of the generation method and available parameter inputs may be found in Appendix A.

5.4 Chosen Assumptions and Parameters of Generated Pseudo-Data (Common to All Scenarios)

The data generation tool was used to construct pseudo-real data sets, which were used to gauge the efficacy of inflation estimation methods under various scenarios. It was considered desirable to retain the majority of the parameters and assumptions as unchanged across all scenarios, with only the inflation parameters varying. This section of the paper summarises these common parameters and assumptions.

Although the parameters chosen were selected with reference to real data, it is stressed that they are entirely artificial in nature. Readers may well have differing views as to the appropriateness of

these parameters. As such, a copy of the generation workbook and worked examples may be provided on request for readers to conduct their own experimentation.

The initial random seed was fixed at the same value for all scenarios outlined in this paper.

Aggregate attritional claims were assumed to follow a lognormal distribution with a mean of 10,000,000 units and a standard deviation of 1,000,000 units.

For each scenario, individual claims frequency was assumed to follow a Poisson distribution with mean of 50. This frequency was selected to provide a reasonable volume of claims in each simulated year of data.

Claims severity was assumed to follow a single parameter Pareto distribution with observation point of 1,000,000 units and alpha parameter of 2.5. These values were chosen on the basis of being “comfortable, round numbers.” However, we note that this choice of alpha parameter will represent a comparatively volatile (high coefficient of variation) distribution.

Frequency trend was typically kept at 100% (i.e. nil negative or positive trend) barring specific scenarios.

Variability in the inflation and development draws (i.e. standard deviations of mixing distributions) were both set at 15%. It was decided to employ these parameters to better approximate “real-world” data. Claims development, for instance, does not consistently follow a “nice” static pattern in real life.

Individual claims were assumed to follow the development pattern shown below. It was adjusted to be fully developed after ten years, for ease of analysis. These are shown in Table 1.

Table 1. Large claim development pattern across development years (DY)

		DY1	DY2	DY3	DY4	DY5	DY6	DY7	DY8	DY9	DY10
Inflation	Incurred Pattern	4%	24%	43%	57%	66%	73%	79%	84%	90%	100%

It should be reinforced that the data used in estimation are artificially generated. Although the stochastic generation method will engender a significant degree of variability – in line with real-world claims – the data are free from non-inflationary, systemic effects, such as change in business mix change in reserving philosophy/practice changes in breadth of cover; and so forth.

6. Scenarios of Analysis

6.1 Overview

Our goal is, on a practical basis, to evaluate the various methods which may be used to estimate the levels of inflation present in historical data. In order to do so, it was decided to apply these techniques to a variety of inflation scenarios of increasing complexity and assess their performance in each. Each scenario adds an incremental element to its predecessor, barring the final scenario.

The data generation tool allows for (additive) inflation loads in economic, social and excess/shock categories; as well as a trend of increasing/decreasing frequency. These inputs were used in setting the various scenarios described below.

6.2 Scenario Descriptions

6.2.1 Scenario A – Constant, Stable Inflation

For this scenario, a constant rate of 4% economic inflation of claim values was assumed for all years. This, in effect, represents the “theoretical ideal” inflation scenario. This assumption is summarised in Table 2.

Table 2. Total inflation values assumed for Scenario A

Calendar year	2014	2015	2016	2017	2018	2019	2020	2021	2022	2023
Inflation	4.0%	4.0%	4.0%	4.0%	4.0%	4.0%	4.0%	4.0%	4.0%	4.0%

6.2.2 Scenario B – Emerging Social Inflation

Scenario B represents the emergence of social inflation. We start with the parameters used in Scenario A (i.e. 4% constant economic inflation in all years) and to this we add an additional 2% social inflation load in origin years 5–10 (inclusive). Thus, total inflation of claim amounts is 4% for years 1 to 4 and 6% for years 5–10 and future years thereafter

This scenario is designed to represent a trend of social inflation which is not initially present, but which manifests itself over time. Parameters for this scenario are shown in Table 3.

Table 3. Total inflation values assumed for Scenario B

Calendar year	2014	2015	2016	2017	2018	2019	2020	2021	2022	2023
Inflation	4.0%	4.0%	4.0%	4.0%	6.0%	6.0%	6.0%	6.0%	6.0%	6.0%

6.2.3 Scenario C – Sudden Shock Inflation

This scenario is identical to Scenario B, but for the addition of a “shock” inflation factor of 6% in year 9 and 4% in year 10. Thus, total inflation of claim amounts is set at 4% for years 1–4; 6% for years 5–8; 12% in year 9; and 10% in year 10. Future years’ (beyond the latest point of the triangle) inflation is then assumed to gradually revert to 6% (i.e. it is assumed that social inflation will persist) five years following the endpoint of the triangle.

Scenario C was designed to represent something of the state in which arguably Casualty classes may have been recently (2021–2023). In other words, inflation was reasonably stable historically; was said to be impacted by rising claim costs due to social inflation in the recent past; and has (in theory) been subject to additional supernormal economic inflation in very recent years.

Inflation parameters for Scenario C are shown in Table 4.

Table 4. Total inflation values assumed for Scenario C

Calendar year	2014	2015	2016	2017	2018	2019	2020	2021	2022	2023
Inflation	4.0%	4.0%	4.0%	4.0%	6.0%	6.0%	6.0%	6.0%	12.0%	10.0%

6.2.4 Scenarios D1 and D2 – Non-Inflationary Frequency Trends

Scenarios D1 and D2 represented both sides of the frequency trend coin. Both used the same inflation parameters as Scenario C. However, in D1, a trend of 5% decreasing mean frequency was assumed and 5% increasing for D2. In other words, claim frequency for scenario D1 was assumed to follow a Poisson distribution with mean of 50 in year one (base parameter), 47.5 in year two, 45.125 in year 3, etc.

These scenarios represent a world where claim costs (i.e. severity) will be rising, but other external factors will be acting to increase or decrease the likelihood of claims occurrence over time. For instance, increased use of parking sensors, maximum speed limit reductions and increased police speed checks are all likely to reduce the frequency of motor claims. Equally, economic downturns and cuts to law enforcement funding are likely to result in an increase in crime rates and associated claim frequencies.

6.2.5 Scenario E – Unknown (to Analyst) Inputs

In each of the previous scenarios, the input inflationary values and frequency trend were known to the analyst performing the estimation work in advance. Arguably this approach reduces the level of objectivity that can be applied in estimation. In other words, if the analyst knows what the answer “should be,” then, when applying judgement, they will naturally be biased thereto.

Accordingly, for the final scenario, one of the working party members set the inflationary and frequency trend inputs (other parameters remained the same as per the other scenarios), whilst a second analyst carried out the estimation work “blindly.”

The chosen parameters were as follows:

- Constant economic inflation of 5.5% in all years
- Social inflation factor of 1% in years 1–4, 2% in years 5–8 and 3% in years 9 and 10
- First inflationary shock which emerges at 3% in year 1, peaks at 4% in year 2 and falls to 2% in year 3, before dissipating
- Second inflationary shock which emerge at 2% in year 9 and rises to 8% in year 10.

Total inflation parameters across all years are therefore as per Table 5.

Table 5. Total inflation values assumed for Scenario E

Calendar year	2014	2015	2016	2017	2018	2019	2020	2021	2022	2023
Inflation	9.5%	10.5%	8.5%	6.5%	7.5%	7.5%	7.5%	7.5%	10.5%	16.5%

In addition to these inflationary parameters, a negative frequency trend of 10% per annum was also applied to the Poisson mean frequency parameter for the large, individual claims.

Lastly, in order to test the robustness of the estimation methods to data volumes, the simulation was produced under decreasing mean frequency values of 50, 40, 30, 20 and 10 claims per annum; with an estimation exercise carried out against each. This “reducing frequency” check was performed for Scenario E only.

7. Inflation Estimation Methods

7.1 Overview

In this section, we provide a description of some of the techniques that may be used to estimate inflation in historical claims data and which have been investigated by the Working Party. Each of these has been well-described in existing literature and so only a brief overview is provided here. Additional descriptions of the methods and our experience in using them may be found in Appendix B.

To our minds, the various techniques explored can be categorised into two groups: inflation estimation methods that may be applied to individual claim estimates and methods applicable to claims in the aggregate (i.e. aggregate claims triangles).

A key recommendation to users employing these methods would be to make heavy use of graphical observation. Graphs of the movement in individual or aggregate claims per period of time, not to mention split by origin year, calendar year or settlement year may highlight trends and features that assist with the selection of the fitting method. It will be useful to aim to standardise such charts across a team or organisation, to facilitate trend observation and identification.

Furthermore, in order to achieve the most accurate result, we observed that a degree of back-fitting to be quite powerful. In other words, first assess the trend in, say, frequency or severity, and

then use this assessment to inform the choice of parameter or index to apply to the historical data that *removes* all trend, with graphical aids used throughout.

7.2 Methods Applicable to Individual Claims

Should granular, individual claims data be available, these methods can prove highly useful through an ability to remove or adjust individual claims, as well as to examine trends in individual percentiles, as opposed to simply trends in the mean. In addition – as will be shown in the analysis of the results, particularly for Scenario D – individual claim methods can more readily distinguish between the frequency and severity trend elements of inflation.

A drawback of these methods, however, is that they do not explicitly account for claims development. In other words, frequency and severity may appear to be trending downwards owing to claims not yet reported or reported claims being insufficiently reserved (as noted above). This shortcoming can be mitigated through performing an inflation study using settled claims data only (as discussed in Appendix B). However, such an approach denies us the information present in open claims.

Three methods have been considered here, each of which is described in additional detail in the appendix and in (Sheaf et al., 2005), namely:

- Large severity trend
- Large frequency trend
- Trend in burning cost to a theoretical layer.

All three of the techniques require the user to, effectively, set a threshold claim size for analysis. The results of the investigation can vary with this threshold selection and so careful consideration should be paid thereto.

These methods are all relatively uncomplicated to apply and so can be readily picked up, employed and tested by the reader. More complex methods not considered in our work include, for instance, the joint-model of longitudinal payments and time-to-settlement for modelling stochastic incurred but not enough reserved (IBNER). This framework was originally developed in the field of biomedical statistics (Rizopoulos, 2021) and attempts have been made recently to adapt it to the field of insurance (Okine et al., 2021). The interested reader is encouraged to explore this developing area further.

7.3 Methods Applicable to Aggregate Claims

The methods described in the subsequent sections are to be applied to claims triangles and tend to be used to estimate inflation on an origin-year basis (compared with settlement year).

For the most part, these techniques were first discussed several decades ago (in the late 1970s and early 1980s, before most of the Working Party had been born), close to the time when hyperinflation was last a serious economic concern. The use of these techniques was thus likely driven by a lack of both the granular data (or easy access to it) and the ubiquitous computational power we have today.

Despite these points, aggregate inflation estimation methods may lend themselves well to attritional claims, where trends are likely to be more stable and tails (at least on an incurred basis) may be relatively short (obviating the settlement versus origin concern to a certain extent). Equally, individual (settled) claims data may not always be available to us in the real world; for instance, if using external data.

The techniques considered under this category are:

- The inflation-adjusted chain ladder (IACL) (Institute and Faculty of Actuaries, 1997)
- The separation method (Taylor, 1977)

- The calendar year development ratio (CYDR) 12–60 method (Lynch & Moore, 2022)

Again, these techniques were investigated for ease of use and adoption. The interested reader might also consider Bayesian chain-ladder Markov Chain Monte Carlo models, such as Wuethrich & Salzmann (2012).

8. Evaluation of Estimation Methods on Chosen Scenarios

8.1 Overview

In this section, we discuss the results of applying the inflation estimation methods mentioned against each of the scenarios considered.

Considering again, briefly, the generated data used in this work, we note that for each scenario ten simulations were run. The estimation methods were applied to each simulation within each scenario. The methods were also applied to the ten simulations in aggregate. Again, we note that each single simulation could represent the large loss experience from a single insurer's class of business. Then, the aggregate of these might represent ten "insurers' worth" of such data that might be available to a reinsurer, broker or consultant.

Accordingly, for each scenario, a relatively large number of inflation estimates are produced (one estimate per simulation per method), with these estimates often taking the form of an index by year.

In each of the sections discussing the results for each scenario, we thus show:

- A scatter plot of the individual estimates per simulation and method (Scenario A only)
- A line plot comparing the "All Sim" (i.e. aggregation of all ten simulations) results for each method
- A final plot comparing the input inflation against the individual simulation estimates (scatter) and "All Sim" estimate (line) for the method we found most accurate for that particular scenario.

By the term, "most accurate" in the third bullet above, we mean the method which we found to most accurately match the selected input inflation. There is a degree of judgement involved here, as the closeness of match will vary with each individual simulation. Our selection of most accurate method reflects our view of the method which most closely approximates the behaviour of the scenario in question, rather than any particularly scoring metric. For scenario E (the "unknown" scenario), this most accurate method was not chosen by the user performing the estimation exercise "blind", but rather when the input parameters were compared against this "blind" application of the methods.

Note that, in Figures 1–16, we observe nine years of estimates. We have simulated ten years of data in each case, meaning there will be nine year-on-year changes to exhibit.

These graphical representations are accompanied by suitable commentary as to any insights gleaned.

9. Results and Commentary by Scenario

9.1 Scenario A – Stable Inflation Across All Years

For this scenario, the burning cost method was found to most closely match the input inflation, although most methods performed comparably well.

Unsurprisingly, all but one method returns a near-identical result at the aggregate simulation level, which matches the input parameter near-exactly. Somewhat surprisingly, the frequency trend method estimate is considerably in excess of this, however. On careful review, we have found




Figure 1. Scenario A – sim-level output comparison.




Figure 2. Scenario A – aggregate simulation output comparison.

this to be a pure anomaly, rather than a calculation error; and attributable to threshold selection and small sample size. Also, it was generally found that the frequency method can overestimate inflation, when inflation levels are low.

Equally, we note that, even in this simplistic scenario, there is a high degree of variation at the individual simulation level (Figure 3). Although our variability in severity is rather large (see inputs), this does highlight the difficulty an individual insurer may have in gauging true inflation (versus process error) using solely their own data. Hereafter (i.e. in subsequent scenarios), the individual simulation chart will not be shown, owing to the noise and difficulty in interpreting it.




Figure 3. Scenario A – selected method (BC) output versus input.




Figure 4. Scenario B – aggregate simulation output comparison.

9.2 Scenario B – As A but with Emerging Social Inflation

Here we see that both the burning cost (Figures 4 and 5) and frequency trend methods (Figure 4) perform reasonably well in spotting the “uptick” in inflation which occurs in 2018 (year five), although the burning cost method appears to be a somewhat closer match. As such, the burning cost method was selected as our chosen, most accurate method in this instance. Again, however, we see considerable variability at the individual simulation level. This will be a result of inherent




Figure 5. Scenario B – selected method (BC) output versus input.




Figure 6. Scenario C – aggregate simulation output comparison.

volatility at the individual simulation level, coupled with the user selecting their “best view” of inflation to remove trend.

It was noted that the severity trend method was highly sensitive to the choice of large loss threshold applied. Although not shown in the figures, at a threshold of 1m units, the method proved a reasonably good match for the input data. However, it produced lower estimates (relative to input) at higher thresholds. Equally, the IACL and separation methods seemingly estimate the original input inflation but miss the upswing.




Figure 7. Scenario C – selected method (frequency trend) output versus input.




Figure 8. Scenario D1 – aggregate simulation output comparison.

9.3 Scenario C – As B, but with Additional Shock Inflation

Scenario C is essentially identical to Scenario B, with an additional inflation uplift in the final year. Again, we see the burning cost and frequency estimation methods perform well, though the burning cost method does seem to overestimate inflation in early years. The frequency trend method has been selected as the most accurate method in this instance. Interestingly, we observe a tighter “spread” of results for the frequency trend method at the individual simulation level. Perhaps most surprisingly is that the severity trend method performs substantially better here than in Scenario B, despite both scenarios being ostensibly near-identical. This is unlikely to be




Figure 9. Scenario D1 – selected method (BC) output versus input.




Figure 10. Scenario D2 – aggregate simulation output comparison.

wholly attributable to simulation error, as, owing to the fixed seed, most stochastic elements in both scenarios should be identical. It is more likely to be attributable to variations between users' judgement; thus highlighting that substantial care and manual, individual input is required in selecting estimates.

9.4 Scenarios D1 and D2 – as C, but with Decreasing and Increasing Mean Frequency

These scenarios are considered together, being essentially both ways of considering the same effect, namely positive and negative frequency trend. As a reminder, these scenarios are, essentially, scenario C but with a negative 5% mean frequency trend per annum in D1 and positive




Figure 11. Scenario D2 – selected method (frequency trend) output versus input.




Figure 12. Scenario E50 – aggregate simulation output comparison.

in D2. Of interest here is the question of how well the various methods can “look through” the frequency effect to ascertain the underlying claim cost trend (i.e. the selected input inflation).

As we add a frequency trend effect, we see the aggregate claim methods perform quite poorly in terms of estimation ability, with frequency and claim cost trend effects becoming inseparable, somewhat, in the aggregate. Equally, the frequency trend method appears to significantly underestimate (cost) inflation in this case, unsurprisingly. Though better performing, even the burning cost approach is not necessarily ideal. It has been selected as the chosen method in this instance.

Paradoxically, the various estimation methods seem to handle a positive frequency trend much better than a negative frequency trend. We would have expected the frequency method to significantly overestimate claims inflation (a sort of amplification effect), but this does not appear to have been the case; with it actually performing best (hence selected, as above) across all methods and the burning cost approach exhibiting said amplification.




Figure 13. Scenario E50 – selected method (severity trend) output versus input.




Figure 14. Scenario E40 – selected method (severity trend) output versus input.

9.5 Scenario E – Unknown (to Analyst) Inputs

As discussed in Section 6, in this Scenario the input inflationary values and frequency trend were unknown to the analyst performing the estimation work. Our intention was to ascertain how accurately an observer with minimal prior knowledge might be able to determine the true, underlying inflation. After the analyst had performed their estimation work, the results were made known to them, for the purposes of producing exhibits and commentary.

9.5.1 Full Claim Set

As per Scenario D1, a negative frequency trend was entered, in addition to the inflation parameters. Just as with Scenario D1, we thus observe the aggregate claim methods perform quite poorly in terms of estimation ability – again, these methods conflate frequency and loss cost trends.

The burning cost method performs well for the early years, though it does seem to underestimate inflation in later years, where it overcompensates for the frequency trend. The severity trend method performs well on an “All” simulation level, slightly underestimating




Figure 15. Scenario E30 – selected method (severity trend) output versus input.




Figure 16. Scenario E20 – selected method (severity trend) output versus input.

inflation for certain years. However, there is a wide range of outputs across the different simulations, owing to the volatility in individual simulations. Severity trend is our selected method in this instance.

It should be noted that, when embarking on this exercise, we had envisaged the severity method to be used as an “initial gauge” to inform selections on other methods. As such, we did not make use of it for determining a varying (by year) inflation index. However, there is nothing to preclude using the method in this manner, similar to the frequency and burning cost methods.

Equally, we note, as mentioned previously, that all of the individual claim methods can be applied “in reverse”. In other words, rather than calculating the trend implied by a given method; we may instead attempt to back-solve for the overall average inflation or inflation index which de-trends the inflated data.

9.5.2 Robustness of Results at Decreasing Claim Counts

Finally, in order to test the robustness of the estimation methods to data volumes, the simulation was produced under decreasing mean frequency values of 40, 30, 20 and 10 claims per annum

(compared with the original 50); with an estimation exercise carried out against each. This “robustness with regard decreasing frequency” analysis was conducted for Scenario E only; as it was envisaged the effect would be similar over all scenarios.

As shown in the results below, the severity method is surprisingly robust at means of 40 and 30 claims (on an “all simulations” basis). However, the accuracy reduces significantly at 20 claims and failed to predict entirely at 10, owing to lack of claims in excess of threshold. The chart showing the results under a frequency assumption of ten claims per annum has thus been omitted.

10. A Note on External Indices

Leaving aside, for now, the question of estimating historical claims inflation using claims data, let us briefly consider the use of benchmark indices as a guide to claims inflation. This is a reasonably common approach across the market. The essence of this is to assume that claims inflation will follow the trend in one or more (i.e. a blend) of available economic indices, such as price inflation, wage inflation, etc.

The approach may be applied with varying degrees of sophistication, as discussed, for instance in the 2022 GIRO workshop *Economic Inflation and Impact on Reserving* (Paul Goodenough and Arti Verma). At a simple level, this may be purely picking an index and adding a fixed load per year (e.g. inflation as per Consumer Price Index (CPI) + 3%). A more thorough approach, however, would be to determine the heads of damage which constitute the claims cohort under investigation; obtain indices which best reflect these heads of damage; and blend the percentage changes in these indices to select a benchmark inflation rate. This reasoning is taken a step further in, for instance (Bohnert, 2015).

The relevance of indices to claims inflation will vary heavily by line of business. First-party motor claims will clearly be closely correlated with an index of used car prices, for instance. However, links between liability classes and available indices are often less apparent; with recent experience showing a degree of decoupling (i.e. there exists some empirical evidence to suggest casualty inflation was not heavily impacted by post-COVID supernormal economic inflation, though the effect may equally simply be delayed and not yet manifested).

Linking claims inflation to an external index may be popular with a variety of stakeholders in an organisation. An index represents a source of “factual” information to which users can point as a benchmark. Underwriters can reference an uptick in wage inflation as a reason to demand higher rates, for instance, whereas making the same argument based on an internal estimation exercise may be more challenging. Even within an organisation, the use of an index can aid in consistency of assumptions between departments (capital, reserving, pricing, etc.)

Further discussion on the use of indices is not in the scope of this paper. Rather, we consider it to warrant a separate investigation, which the Working Party aspire to tackle in 2025. However, we would like to draw attention to one point; namely that deriving an inflation assumption from economic indices is, essentially, a derivation of settlement-year basis inflation. This should not necessarily be applied as an origin-year trend without adjustment or, at least, consideration as to its appropriateness. This point will be discussed further in Section 11.

11. Additional Challenges in Inflation Estimation

11.1 Gearing of Gross Inflation into Reinsurance or Excess Layers

As discussed in detail in Sheaf et al. (2005), it is a reasonably well-known phenomenon that an annual level of cost inflation of, say, $x\%$ in gross claims will typically result in $y \geq x\%$ to a corresponding reinsurance or excess layer. Put concisely, this will be due to an additional frequency effect, whereby claims previously below the retention of the layer are “pushed into it” by inflation.




Figure 17. Comparison of gearing effects.

A potential method to mitigate the impact from gearing could involve de-indexing the layer historically. However, this methodology becomes somewhat circular with the de-indexing values selected directly influencing the resultant estimates of historic inflation. Given this outcome, the layers have not been de-indexed for the purposes of the inflation estimation methods shown in this paper.

Accordingly, when applying the burning cost trend method to estimate inflation in our examples, we are explicitly estimating the cost inflation in the applied, theoretical reinsurance layer. This inflation parameter must then be converted back into a gross estimate. Where claims severity follows a single-parameter Pareto distribution, this conversion is entirely formulaic and independent of the layer in question.

However, severity can follow (or to be more accurate, be best approximated by) a wide variety of probability distributions. The level of gearing can vary quite dramatically between distributions and even for different reinsurance structures applied.

Figure 17 illustrates this, where we use 10,000 gross claim simulations to compare the gross versus geared (to reinsurance layer) inflation. This has been carried out under the originally specified Pareto distribution used throughout the rest of the estimation work and with the same reinsurance layer (5 million units of limit in excess of 2 million units attachment) applied, as well as for:

- A LogNormal distribution with the same mean and variance as said Pareto distribution and same reinsurance layer
- That same LogNormal distribution, but under the application of a 5 million in excess of 5 million (units), theoretical reinsurance layer.

What may be seen is that the level of gearing differs markedly across the three cases. For equivalent mean, volatility and layer parameters, the Lognormal distribution produces a substantially smaller level of gearing. However, when the attachment point of the theoretical layer is increased, the level of gearing increases dramatically. With this said, we do observe a power curve relationship between the gross and geared inflation in all cases.

Given this observation, the usefulness of the burning cost trend method for ascertaining the gross inflation level is somewhat called into question. For a known level of gross inflation, a simulation approach can easily be used to estimate the level of geared/excess inflation (as the severity distribution may be ascertained). However, the converse does not hold true quite as easily, as we cannot *know* the gross severity distribution without first on-levelling the gross claims . . . for which we need the gross inflation parameter.

A user could, conceivably, produce many instances of the above simulation for a variety of different distributions, parameters and reinsurance layers and then use these to “back-solve” a gross inflation parameter from any geared estimate derived empirically from the burning cost method. However, this strikes us as a highly time-consuming exercise with questionable accuracy. Rather, we suggest that the burning cost trend method ought to be used cautiously, for additional guidance and under the “neat” single-parameter Pareto method, instead of being overly relied on for the selected inflation level.

Note that consideration of gearing effects will implicitly also need to be given if the user wishes to separately estimate inflation for different claim categories – i.e. attritional versus large claims.

12. Settlement vs. Origin-Year Inflation

12.1 Overview

In *applying* any inflation estimates – i.e. on-levelling – we will invariably be required to do so on an origin-year basis; simply because this is the basis on which all insurance operations are conducted. (Re)insurance policies are priced on the basis of inception date and reserves are estimated on the basis of accident/underwriting year; as are capital requirements. Simply put, there is rarely a *need* to explicitly think about the quantum and distribution of claims that will *settle* in a given future year.

However, as we have alluded to, inflation is often optimally estimated on a settlement year (i.e. year in which claim is paid) basis. External indices effectively align with claims settlements, as mentioned, and settled claims have the advantage of being largely “locked down”. In other words, for any given historical year, we have near certainty as to the volume and quantum of claims settled in that year.

This temporal dichotomy between period of estimation and period of application leads to a challenge, for which we consider there to be three possible mitigating approaches:

1. By and large ignore the settlement-year/origin-year dichotomy
2. Explicitly estimate inflation on a settlement-year basis and allocate to origin year
3. Estimate inflation on an origin-year basis, but with an adjustment to the original data to approximate settled values.

Each of these approaches will now be discussed in detail.

12.2 Approach 1 – Minimal Change

This approach is less worrisome than it may initially sound. For short-tailed lines such as property and motor own damage, case reserves are likely to be reasonably reflective of final settlement values (with the exception, perhaps, of particularly large catastrophe events and/or particularly complex risk claims) and there will be minimal development in both incurred values (i.e. incurred but not enough reserved – IBNER) and claim frequency (i.e. incurred but not yet reported – IBNYR) once an origin year has ended (slightly longer on underwriting year basis than accident year).

Accordingly, for such short-tailed classes, a combination of open and closed claims may potentially be used in estimation, with said estimation being conducted on an origin-year basis.




Figure 18. Example payment pattern (10 year).

The most recent (open) origin year perhaps ought perhaps to be excluded, in addition, to avoid the distorting effect of development here; particularly if working on an underwriting year basis, where the most recent year may not be fully earned.

Unfortunately, this approximation approach is likely only to be valid in quite short-tailed, first-party lines of business. For such lines, difficulty in *estimating* inflation tends to be less of an issue anyway, as the drivers of claim costs tend to be reasonably evident. For instance, if we observe an increase in building material costs, it may reasonably be expected that property claims in respect of rebuild costs will increase also.

Equally, the approximation will rely on the implicit assumption of a historically stable portfolio in terms of mix, terms and conditions, reserving philosophy, etc. Reserving philosophy is of particular importance, as a deceleration of reporting or weakening of case reserving strength (or postal strike as was put to the authors in their student days) can manifest itself as a negative trend (and vice versa). Again, though, this will be less of a concern when the tail is short.

12.3 Approach 2 – Allocation of Settlement-Year Inflation to Origin Year

Here, we make use of settled claims only to carry out the inflation estimation exercise. Following this, we apply the logic that a claim in origin year x will settle in any of years x to $x+n$ (where n represents 100% development of the applicable payment pattern). As such, we can consider the inflation to be applied in on-levelling claims from origin year x to be a weighted average of the inflation in settlement years x to $x+n$, with weights determined by the payment pattern.

This approach may appear elegant on initial reading but throws up two challenges. Firstly, making use of settled claims only (unless we're simply using an index) ignores any information in open claims. In general insurance and, in particular, the London Market, dearth of data is often bemoaned and so actively disregarding data feels distinctly discomfiting.

Secondly, the approach will require us to possess an estimate of settlement-year inflation in *future* settlement years – which is fundamentally unknowable with certainty. For settlement years in the near future, users may apply judgement to make, essentially, an educated guess. However, for casualty lines of business, payment patterns may stretch two decades or more and the user will likely need to revert to some form of long-run average settlement-year inflation assumption. Accordingly, using allocated (to origin year) settlement-year inflation for on-levelling may be thought to engender a considerable degree of uncertainty in the process.

We provide a brief illustration of this point in Figure 18. Consider a class of business with a ten-year payment pattern for a given origin-year triangle. The pattern implies a mean time to payment of circa 6.6 years, so the tail is not overly extreme.




Figure 19. Origin-year inflation attributable to past/future settlement-year inflation.

Let us imagine that we are performing an on-levelling exercise as at 31/12/2023 and have obtained settlement-year inflation estimates for the 2023 and prior settlement years. Figure 19 then illustrates the proportion of inflation in each origin year attributable to historical settlement-year inflation (i.e. “known” estimates) versus prospective settlement-year inflation (i.e. “unknown” and “unknowable” estimates).

As can be seen in Figure 19, for each of the five most recent origin years, no more than 25% of the inflation assumption used to on-level data can be attributed to known inflation estimates. This result is entirely independent of the inflation estimates themselves and is purely a function of the payment pattern.

Given the significant influence of the estimates of future settlement-year inflation in on-levelling, the reader may (at first glance) question the merits in estimating historical settlement-year inflation altogether. We contend that this exercise is still highly worthwhile, not least because the estimates will be directly used but also because historical estimates of settlement-year inflation will be the best source of insight for informing the selection of future estimates.

12.4 Approach 3 – Adjust/Transform Origin-Year Data and Estimate Inflation on Origin-Year Basis

In essence, the thinking here will be to apply a level of development to the original data on an origin-year basis, before implementing inflation estimation techniques. Two approaches might be considered here – flipsides of the same coin. Namely: assume all development in claims is in respect of IBNYR, not IBNER and vice versa.

In the case of the former assumption, we are explicitly assuming that case reserves for open claims are wholly reflective of final settlement values and that all “Incurred But Not Reported” (IBNR) liability is in respect of claims yet to be reported. Thus, we will need to estimate the number of claims “yet to be reported” in each origin year and generate a severity for each (via a stochastic process).

In the case of the latter assumption, we will instead be assuming that all claims are known (a not necessarily unreasonable assumption for claims made business, for instance), but that the final, settlement values of these claims is not. Here we will need to apply an IBNER development method to open claims to determine settlement values. Such a method is described in (Parodi, 2023).

In either case, we will, essentially obtain a listing of ultimate, settled claim estimates by origin year, on which individual claim estimation methods may be applied to determine inflation estimates by said origin year. The Working Party have not investigated these approaches in any

depth (owing to considerations around sharing “real” data amongst the working party). However, we do have concerns that these methods:

- May result in a certain level of circularity
- Are cumbersome and opaque to perform
- Will rely heavily on expert judgement
- May double count inflation in applying development.

13. Closing Remarks

Let us return to our opening statement:

“Unfortunately, claims inflation is notoriously difficult to measure with any degree of certainty”

In our exploration of claims inflation estimation, we have certainly come to agree with this statement. Outside of highly simplistic cases (e.g. flat inflation), the various estimation methods require considerable care and judgement in their employment. Equally, a considerable volume of data is required to separate the effects of inflation from random variability or process error (although admittedly, our parameterised volatility was large).

In general, we observed that methods applicable to individual claims performed more successfully than those applicable to claims triangles, owing to being better able to separate frequency and severity effects. However, these approaches will be time-consuming and work best when a degree of validation is employed, e.g. verifying removal of trend post the application of a selected index estimate. Equally, they are reliant on availability of individual claim data, with additional concerns around the origin- versus settlement-year question.

Furthermore, the usefulness of the burning cost trend method is somewhat called into question, owing to the difficulty and uncertainty in “reversing out” the gearing effect to determine the gross inflation estimate. The frequency trend method also seems to encounter issues in that it may overestimate inflation when underlying inflation is low (or even, as can occasionally transpire, negative).

Both the frequency trend and severity trend methods were highly sensitive to the choice of threshold parameter. The inference to be drawn here actually pertains to communication of result, rather than validity of method. In other words, after conducting an inflation estimation exercise with a method involving a threshold selection, the analyst should carefully communicate this to any users – in essence that the estimated parameters should, ideally, be used for claims above the selected threshold.

It firmly remains our belief that inflation estimation is a valid and worthwhile exercise to perform, as its use in on-levelling has a bearing on all aspects of actuarial activities in general insurance. However, we equally believe that it is an exercise laden with judgement, is difficult to automate and thus requires a degree of stakeholder engagement and agreement.

13.1 Future Research

Having considered the question of “*what is inflation?*,” the working party will continue to investigate how inflation impacts our work as actuaries in general insurance. As stated at the outset of this paper, our goal is to produce a variety of useful, practical and digestible guidance reports on the subject.

We have noted already, for instance, that external indices remain one of the most popular, communicable and consensus-building sources of claims inflation. We are keen to explore questions as to how best to map and blend indices to be most applicable to a claim set or class of business of interest. In addition, noting the issues around settlement versus origin-year inflation, *prediction* or forecast of indices becomes of considerable interest.

At a more practical level, many of us will have observed the challenges faced in estimating claims reserves in the latter half of 2021 and well into 2022. The trinity of chain ladder, Bornhuetter-Ferguson and *a priori* methods remain heavily used to set reserves across the actuarial world, but were thrown into question during this period. We are thus keen to investigate the challenge as to how best to conduct a reserving exercise during a time of sudden inflationary change; as well as the perils of “getting it wrong”.

These two topics of using indices to inform inflation estimates and how best to set reserves during periods of inflationary change are likely to form the focal point of our next areas of research. However, the authors are extremely welcoming of feedback as to what would be of benefit to actuaries in terms of areas of future research; as well as to any comments readers may have on the contents of this paper.

Acknowledgements. The working party are immensely grateful to the following individuals for their help in producing this paper and the analysis which underpinned it:

- Melanie McDowell of Marcuson
- Vignesh Balaji of Marcuson
- José Gómez Mena of Marcuson
- Steffan Xenios of Hymans Robertson
- Siddhant Chopra of Hymans Robertson
- Megan Clarke of Hymans Robertson
- Ryan Farnes of Hiscox

In addition, we are most grateful to the following individuals for their work in reviewing this paper:

- Richard Rodriguez of Guy Carpenter
- Yuriy Krvavych of Guy Carpenter.

Disclaimer. The views expressed in this publication are those of invited contributors and not necessarily those of the Institute and Faculty of Actuaries. The Institute and Faculty of Actuaries do not endorse any of the views stated, nor any claims or representations made in this publication and accept no responsibility or liability to any person for loss or damage suffered as a consequence of their placing reliance upon any view, claim or representation made in this publication. The information and expressions of opinion contained in this publication are not intended to be a comprehensive study, nor to provide actuarial advice or advice of any nature and should not be treated as a substitute for specific advice concerning individual situations. On no account may any part of this publication be reproduced without the written permission of the Institute and Faculty of Actuaries.

References

- Avanzi, B. (2001). SynthETIC: An individual insurance claim simulator with feature. *Insurance: Mathematics and Economics*, **100**, 296–308.
- Bohnert, A. (2015). Assessing inflation risk in non-life insurance. *Insurance: Mathematics and Economics*, **66**, 86–96.
- Cairns, M., Lu, A., & Shah, N. (2022). We Don't Have to Worry About Inflation – We're Commercial Lines. *Casualty Actuarial Society*. (2024). CS-8 Double inflation (Micro + Social) – What Can We Do? Boston, Casualty Actuarial Society Reinsurance Seminar.
- Faculty and Institute of Actuaries. (1997). Claims Reserving Manual.
- Flower, M., et al. (2006). Reinsurance Matters. Institute and Faculty of Actuaries.
- Gabriella, A., & Wuethrich, M. (2018). An individual claims history simulation machine. *Risks*, **6**, 1–32.
- Institute and Faculty of Actuaries. (1997). Claims Reserving Manual.
- Lloyd's of London. (2022a). *Claims inflation – What are Lloyd's doing?*. Liverpool: Institute and Faculty of Actuaries.
- Lloyd's of London. (2022b). *Reserving guidance – Allowing for inflation*, London: Lloyd's of London.
- Lynch, J., & Moore, D. (2022). *Social inflation and loss development*. Casualty Actuarial Society.
- Okine, N.-A., & Wuethrich, M. (2021). Joint model prediction and application to individual-loss reserving. *ASTIN Bulletin: The Journal of the IAA*, **52**, 91–96.
- Parodi, P. (2023). *Pricing in general insurance*. Chapman and Hall.
- PRA. (2023). Bank of England. [Online]. <https://www.bankofengland.co.uk/-/media/boe/files/prudential-regulation/letter/2023/insurance-supervision-2023-priorities.pdf>
- PRA. (2024). [Online]. <https://www.bankofengland.co.uk/-/media/boe/files/prudential-regulation/letter/2024/insurance-supervision-2024-priorities.pdf>

- Richards, W. F. (1980). *Evaluating the impact of inflation on loss reserves*. Casualty Actuarial Society.
- Rizopoulos, D. (2021). *Joint modeling of longitudinal and time-to-event data with applications in R*. Erasmus University Medical Centre.
- Sheaf, S., Shuster, S., & Forster, W. (2005). Claims Inflation Uses and Abuses. Insitute and Faculty of Actuaries.
- Taylor, G. (1977). Separation of inflation and other effects from the distribution of non-life insurance claims. *ASTIN Bulletin: The Journal of the IAA*, 9, 219–230.
- Thomson, W. (1958). A modified congruence method of generating pseudo-random numbers. *The Computer Journal*, 1, 83.
- Wuethrich, M., & Salzmänn, R. (2012). Modelling accounting year dependencies in runoff triangles. *European Actuarial Journal*, 2, 227–242.

Appendix A. Description of Claims Generation Tool

A.1 Overview

In this Appendix, the Excel-based tool used to generate the pseudo loss data used in the analysis is described. Some sections of the code are included where relevant to the discussion, but the full code is not included within this paper for sake of brevity.

The tool (which includes the VBA code) is available on request.

A.2 Rationale for Constructing a Tool

There are multiple options for generating pseudo loss data to investigate claims inflation techniques from a variety of software packages. The goal of this paper is to ensure the work presented can be replicated by all interested readers without the requirement for specialised software licences.

The tool has been built in Excel/VBA as this was a common medium that all within the Working Group used. The tool has been built using as few internal Excel functions as possible to enhance portability of the approaches and analysis to other scripting languages, such as Python or R.

A.3 Random Seed

In addition to the requirement of an application available to all in the Working Party, the tool required the ability to generate the same set of outputs for each user. This was achieved by including an input for the random seed to be used for generation of the random numbers used in the simulations.

The generation of random numbers is a complex task, and many applications and coding languages currently use more complicated methods than used in the tool, which is based on a Linear Congruential Generator (LCG) (Thomson, 1958). The parameters used in the tool are based on those used by Microsoft historically.

A fully defined LCG has been used in the interests of portability.

Function rand_num(RdSeed)

a = 214,013

c = 2531,011

m = 2 ^ 32

*rd_num = (a * RdSeed + c)*

int_div = Int(rd_num/m)

*rand_num = rd_num - int_div * m*

End Function

A.4 Use of Excel/VBA as Simulation Tool

Whilst Excel is incredibly flexible and allows the provision of detailed information, its use as a tool to simulate a large volume of random variables can be a slow process.

To assist with the timeframe required to generate the results required for this paper, two methods were employed.

Firstly, the number of simulations was set at 10, with a mean frequency of 50 individual claims, to ensure a range of results that did not overwhelm the number of calculations required. Secondly, the distributions of the aggregate attritional claims and the frequency and severity of the individual large claims are calculated as empirical tables at the outset of each set of

simulations. These tables are then referenced during each simulation as the derivation of these tables is only required once per set of simulations.

A.5 Generation of Pseudo Claims

The claims generated are split between the aggregated attritional claims and individual large claims.

A.5.1 Attritional Claims

These are modelled as annual totals for each simulation, therefore there is only one entry in the tables for a given simulation. The value for these claims is generated using the random number generated and the empirical distribution table, based on user inputs.

The output table includes a value for the total attritional claims amount with no inflation applied, as well as a total value following application of the input development pattern and inflationary impacts from each development year.

Where future development extends beyond the last year of claims inflation input, the last value for claims inflation is used.

A.5.2 Large Individual Claims

These are modelled in a similar manner to the attritional claims, with an intermediate step that determines how many claims should be generated within the simulated year. Once this frequency is known, each claim is generated within a given year within a simulation. The presentation of the results from the simulation for individual claims follows the attritional claims described above, with a total uninflated amount, a total amount after inflation, and a cumulative amount in each development year.

Once generated, the large individual claims are aggregated to year and simulation number level, to allow aggregate claims inflation methods to also be applied.

Appendix B. Additional Notes on Employing the Estimation Methods

B.1 Overview of Section

In this Appendix, additional commentary will be provided on the various claims inflation estimation methods investigated in this work. The commentary will be divided between general observations and explanations of the methods, as well as specific observations and usage context from our employment of these methods in the various inflationary scenarios considered. For completeness, the interested reader is directed to the sources listed in the bibliography.

B.2 Methods Applicable to Individual Claims

B.2.1 Large Severity Trend

Simply put, this approach involves the selection of a fixed (per year) threshold, followed by the calculation of the average and varying percentiles of claim amounts for claim amounts in excess of this threshold in each year. The average trend over a number of years (based on average or percentile severity) can then be estimated using an exponential approximation (e.g. “Logest” function in Excel) or simply by comparing the first and final data points (e.g. comparing 25th percentile severity in years 1 and 10) to compute a simple average trend. It should be noted, however, that these two approaches to calculating the overall period-average trend can lead to quite differing results. The “Logest” approach may be preferred, for instance, as it better deals with distorting effects from (for example) lack of data in the most recent year.

Perhaps more usefully, graphical observation of the average and percentile severities over time can aid in determining an inflation index (as opposed to single average value) over the period of interest; as well as identifying anomalous claims for removal/adjustment. This can be a highly time-consuming and labour-intensive exercise, however. Fundamentally, there will be a trade-off between the need to obtain a detailed, granular index and the time/resource constraints imposed on the analyst. Although judgement and graphical analysis may be preferred, there will be many times when a general view of overall period trend may be sufficient.

In addition, severity trends will be absent from “distorting” frequency effects (e.g. rising frequency due to increasing propensity to claim), which may obscure the “true” view of claims inflation (i.e. trend in claims costs).

Although we have applied severity analysis techniques on individual claims, an alternative approach might be to examine the average claims cost triangle – i.e. the triangle derived from dividing an incurred claims triangle (or paid triangle) by its corresponding claim count triangle. This approach would consider inflation on an origin (as opposed to settlement) year basis.

Such methods come with their own benefits and limitations. By using the average cost of claims, an element of smoothing of the information is included, removing some of the noise inherent in these data sets. However, the more recent elements of the triangle are likely to be the more valuable in any study, but these data points are likely to be less developed and fewer in number.

Additional complications can be employed to address this shortcoming by estimating the underlying development patterns of claims and claims count over time (either from the claims triangle itself or a suitable benchmark) to investigate the trend in average claims of the estimated ultimate claims position. This approach opens itself up to potential distortions in results due to outdated or incorrect development patterns.

For this method, severity was considered for claims in excess of 1.5 m units. It should be noted that the estimation exercise can be quite sensitive to the selected threshold and careful consideration and testing should be given to its selection.

In each year, the mean severity was calculated for claims in excess of this threshold; as well as the severity at the 25th, 50th, 75th, 90th, 95th, 97.5th and 99th percentiles. The long-run trend in severity across the period for each of these metrics was then determined and an overall-period inflation estimate selected. It was generally found, however, that the trends in severity in higher percentiles tended to be overly “noisy” and volatile for estimation purposes.

In addition to the pure inflation calculation, charts were plotted of the severity metrics and changes in these metrics over time. These were used to inform the selection of the inflation indices (as opposed to period-average values) for the IACL, frequency trend and burning cost trend metrics.

B.2.2 Large Frequency Trend

As with severity trend, this is not a particularly difficult technique to employ. In essence, this approach requires the user to determine the count of claims in excess of a chosen threshold per unit of exposure (which itself ought to be on-levelled) and determine the overall trend or trend index. By examining the trend in claims in excess of a threshold, we will be capturing both severity and frequency effects (lower claims “pushed” above threshold via severity change).

A key pitfall of this approach is that, even more than severity trend, the choice of analysis threshold is crucial. If set too low, we can easily underestimate the “true” level of claims severity inflation. If set too high, we risk a data set that is too volatile and small to generate robust results. Equally, this approach does not separate severity from frequency trends.

As per the severity trend method, here we considered claims in excess of 1.5m units for the purpose of inflation estimation. As per the IACL, a whole-period trend was estimated initially (as described in Section 8) and this was then used to estimate the inflation index in each year. Graphical inspection – i.e. comparing untrended, trended at period average, and index-trended claims – was used to assist in selecting this index.

As with the IACL, the determination of the index per simulation was quite a manual exercise. Equally, as per the severity trending method, the estimates were found to be quite sensitive to the selected threshold. Specific care needs to be taken in that if the selected threshold is set too “low”, then there becomes no distinction between inflation estimates after a point. In other words, the frequency estimation method would show inflation estimates of 50% per annum and 5000% per annum as equally “valid”.

B.2.3 Trend in Burning Cost to a Theoretical Layer

Sheaf *et al.* (2005) present a mathematical formulation of this technique, which the reader is encouraged to peruse. For our work, however, we approached the technique in a relatively straightforward manner. Again, working as we were with what we interpreted as settled individual claims (stochastically generated, inflated ultimate claims), our approach was simply to define a theoretical reinsurance/excess layer; calculate the loss to this layer for each claim in each settlement year; and aggregate these layer losses within each settlement year (effectively assuming unlimited horizontal coverage).

The resulting trend in burning cost could then be assessed and both period-average and index trend values estimated. As with all approaches, graphical investigation was key in selecting the estimates.

Here, the reference layer was set as 5 million in excess of 2 million units. Largely speaking, any nuances of approach here were similar to those applied in the frequency trend estimation method.

This approach, being, in essence, severity and frequency (in excess of a given threshold trend) was generally found to be the most accurate estimation method. Care again needs to be taken with regards to the selection of the layer parameters.

It should also be noted that the inflation estimate obtained through analysing the trend here is the *inflation of losses to the reinsurance layer* and not the *gross* inflation parameter sought. However, knowing that the losses generated follow a single-parameter Pareto distribution, the gross inflation parameter is easily derived as follows (Sheaf *et al.*, 2005).

$$i_{\text{gross}} = (1 + i_{\text{layer}})^{\frac{1}{\alpha}} - 1$$

As mentioned above, when the underlying severity distribution is unknown, this method becomes significantly more difficult to employ with any real certainty.

B.2.4 Methods Applicable to Aggregate Claims

The Inflation-Adjusted Chain Ladder

Many actuaries in the UK will recall the IACL method from their professional examinations. Despite it being reasonably widely known, its employment (prior to the COVID-19 pandemic, at least) was not widespread.

Discussed in (Institute and Faculty of Actuaries, 1997) the IACL may be considered as an enhancement to the standard Chain Ladder (CL) loss reserving technique. The basic CL method implicitly assumes that future inflation will (on average) be equal to that observed historically. The IACL, by contrast, effectively seeks to “strip out” historical inflation present in a claims triangle (be it paid or incurred) before explicitly applying prospective inflation assumptions to determine ultimate, inflation-reflective claims.

The latter, prospective part of the IACL is not within the scope of this particular paper (though will be considered when the working party moves on to explicitly consider questions of reserving). However, the act of “stripping out” historical inflation is of particular interest here as it lends itself to an *estimation* of this historical inflation.

The method of estimation essentially requires the user to *deflate* (top left corner by zero years; next diagonal by one year and so forth) each point in the claims triangle by an index or average inflation parameter, such that the CL ultimate estimates of this deflated triangle are entirely free of trend. In other words, the “correct” historical inflation estimate will lead to stability of estimated ultimate positions in the deflated triangle.

Unfortunately, if there exist other sources of trend in the data (e.g. frequency trend), then these will need to be stripped out separately, to correctly estimate the inflation effect. Equally, if using an incurred triangle, then changes in reserving philosophy may also obfuscate inflationary effects. Stronger case reserving over time may well manifest itself as an inflationary effect if not explicitly dealt with.

In implementing the IACL approach to estimating claims inflation, we first formulaically determined the “whole period” parameter that would result in a removal of all trend in ultimate. This was then used to inform the selection of an index of claims inflation by year.

This latter stage was reasonably manual and informed by judgement. In essence, a combination of the period-average estimate and graphical scrutiny was used to ensure the appropriate index, which removed the trend.

The Separation Method

This method is first described in (Taylor, 1977) and was revisited at GIRO 2022 in a workshop session entitled “*We don’t have to worry about inflation. We’re commercial lines*” (Cairns et al., 2022). This approach bears some similarities to the IACL in that it focusses on calendar years (diagonals) within a triangle, rather than origin years (which it requires to be standardised). In essence, this approach approximates a two-factor model and reads off a calendar year index from the fitted parameters.

The Working Party found this method to be reasonably adequate at determining overall average annual inflation within a given time-period, but of limited use in determining the more granular index (i.e. inflation in each year) for that period.

This method was not found to be ideal in terms of determining a varying inflation index by year, as opposed to a single, period-average inflation parameter. As such, when employing this method, the inflation estimate will be constant across all years analysed. Our approach in utilising this method was also reasonably formulaic, without significant application of judgement. It is considered to have performed relatively poorly, in general.

The Calendar Year Development Ratio (CYDR) 12–60 Method

This is a comparatively recently publicised inflation estimation method, described in (Lynch & Moore, 2022) and also discussed in the “*We don’t have to worry about inflation. We’re commercial lines*” GIRO 2022 workshop (Cairns et al., 2022). Simply put, this method again considers diagonals (calendar years) of the claims triangle. The method requires the user to calculate the product of the 12–24 months (i.e. years 1 to 2), 24–26, 36–48 and 48–60 link ratios along a given diagonal of the triangle. This is known as the CYDR 12–60 factor.

Increases in this factor over time provides evidence of increasing inflation. However, this method – although meritorious in *giving an indication or signal as to* inflationary volatility or highlighting changes in underlying claims inflation – does not provide the means to estimate what this inflation is. As such, though it is useful in determining an overall sense of trend and spikes, it is not overly informative. Other uses of the method may be to check if data has been sufficiently detrended (e.g. if IACL has been applied correctly).

This method does not lead to the production of actual inflation estimates (as far as the working party could make out), as mentioned. Accordingly, it does not appear in any results, although it was used to inform judgements.