Institute and Faculty of Actuaries

Reserving for volatile inflation

A practitioner's guide to setting best estimate reserves during periods of inflation volatility

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

*ocriodac@tcd.ie

30 June 2025

Contents

Impetus and Motivation

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 many other nations exhibited a sharp increase. Price inflation reached its peak in 2022. Though it has since reduced significantly, an increasingly volatile geopolitical climate and restrictions on free trade are seeing a reversal of this trend.

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)* and it remained a supervisory priority into 2024 *(PRA, 2024)*. Although inflation is no longer noted as a supervisory priority for 2025 *(PRA, 2025)*, it receives mention nonetheless. In particular, the PRA caution that *"the full extent of claims inflation in long-tailed lines may not yet be visible."*

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, from the perspective of actuarial work in general insurance.

By way of a side note: the career experiences of the Working Party members have been predominantly focussed in the London Market, albeit across a wide spectrum of functions and organisations in this sphere. Although we consciously strive to produce guidance of wide application, this London Market focus may subconsciously colour our thinking.

Introduction

Overview

As mentioned, the working party was constituted at the beginning of 2023, with the aim of producing practical guidance on the management of inflation uncertainty. It was recognised from the outset that inflation volatility or a period of inflationary instability will have impacts across the spectrum of general insurance actuarial activity. However, before addressing questions as to how to best account for inflationary change or volatility in provisioning, pricing or capital-setting; it was first deemed necessary to ascertain precisely what levels of inflation for which to allow.

This conundrum is considered in two facets, namely:

1. What level of inflation exists in the historical data for the claim cohort of interest; and
2. What is the link between this claim inflation and observable levels of external inflation (as measured by relevant indices)?

The first question was largely addressed in the working party's 2023 research, presented at the 2023 UK GIRO conference and to be found in our paper on that subject *(Creedon, Bargate, Lenney, Schofield, & Stock, 2024)* as well as the additional sources noted therein. By and large, this topic is considered to have been suitably covered for our purposes.

Considered in concert with the first question, a thorough consideration of the second will enable practitioners to identify the core drivers of claims inflation for a given cohort, as can be measured externally (i.e. through indices). This will enable the use of more accurate inflation assumptions in actuarial practice. The working party began to undertake research into this area in 2024 and expects to advance it in 2025.

In parallel with the research undertaken into linking externally measured inflation with internally estimated claims inflation, however, the Working Party also commenced its assessment of the practical impacts of inflation volatility in 2024. In particular, consideration was given to the area of provisioning or reserving; under the working assumption that the reserving actuary is reasonably confident in their views of the inflation parameters. This then forms the question of interest for this paper, namely:

During a period of inflation volatility, how can a practitioner most accurately estimate provisions for claim liabilities?

Outline of Report

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 the first section we outline our understanding of the term "claims inflation". This topic was covered in our previous paper *(Creedon, Bargate, Lenney, Schofield, & Stock, 2024)* and, accordingly, will be covered reasonably perfunctorily here.

Next, we discuss the data used in all of the estimation research we conducted; as well as the assumptions and approach used to generate those data. Full information on the data generation approach is deferred to the appendix. We also discuss how we used these data to create a variety of 'scenarios' under which claims provisions were estimated. The data and scenarios used for this research are precisely the same as those used in our previous paper on claims inflation estimation *(Creedon, Bargate, Lenney, Schofield, & Stock, 2024)*. This continuity has been advantageous to the Working Party and, it is hoped, will be also to the reader.

Following on from this point, we will describe a variety of potential approaches which may be employed in setting or adjusting claims reserves during periods of inflationary volatility. By and large, these are the methods put forward by the Corporation of Lloyd's in their various guidance issuances of 2022, such as their GIRO presentation of that year *(Lloyd's of London, 2022)* and their Reserving Guidance paper *(Lloyd's of London, 2022)*.

As with our previous paper on inflation estimation, each of the approaches described will be evaluated using our generated scenario data (using the same scenarios as in that previous paper). This evaluation will include a comparison of 'known' (i.e. as generated) ultimate and incurred but not yet reported (IBNR) positions against best estimates as-at the latest position of the triangle.

Although the approaches described and evaluated in these sections of the paper were those most commonly employed in the London Market in recent years (with the Cashflow Uplift method being highlighted by Lloyd's as a "good approach", *(Lloyd's of London, 2022)*), other approaches to general insurance claim provisioning assume greater prominence elsewhere. In particular, as noted in the Casualty Actuarial Society's June 2024 reinsurance conference *(Casualty Actuarial Society, 2024)*, there are a variety of reserving methods whose status had been relatively ubiquitous in Australia (though more rarely employed in the UK and USA) considerably prior to the post-pandemic supernormal inflation and which require the user to explicitly consider inflation, both past and future. A section of the paper will provide a concise discussion of these methods, as well as a brief evaluation vs. the Cashflow Uplift approach.

The penultimate section of the paper considers some of the ancillary issues which arose during the provisioning work; their practical implications; and ways to manage them.

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

Definition of Claims Inflation

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, 2022)*, 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 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 a term which has come into use relatively recently (within the last decade or less) when discussing claims inflation. 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 can be considered claims inflation (when considering a cohort of claims in their aggregate); just as an increase in claims costs with no change in incidence rate will be also.

Reserving is, in principle, an exercise in setting an overall, aggregate level of provisions to be held in order to meet claims' liabilities. As such, inflationary effects will, too, often be considered in aggregate. However, such simplifications may lead to a degree of distortion in provision setting, as will be discussed in our section on Ancillary Considerations.

Words of Caution

The caveat put forth in our previous paper *(Creedon, Bargate, Lenney, Schofield, & Stock, 2024)*, though hackneyed, applies equally here; namely that change is a constant. Insurance is, in essence, a business which reacts to the wider environment, with both that environment and the business of insurance itself existing in perpetual states of flux.

Within an insurance entity, a cohort of business will alter over time, even from the action of internal stimuli alone: through changes in mix, terms and conditions, technology, case reserving philosophy, etc. Externally, the environment in which that cohort operates will, equally, not remain static.

One of the subjects of this paper has been to evaluate reserving approaches on a number of pseudo-data scenarios. In providing these worked examples, 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.

In practice, however, inflation will be one of but many sources of uncertainty in estimating claims' provisions. Separate excogitation of frequency and severity trend is certainly considered to be of significant import in a reserving exercise. Equally, the authors hold to the maxim that collaboration and communication across the breadth of an organisation are pivotal. By way of an aside; this principle on the necessity of internal collaboration (and the dangers presented by internal dissociation) formed a key focus of the author's lecture on actuarial professionalism in May 2025 *(Creedon, Actuarial Professionalism, 2025)*.

At its heart, provisioning is an exercise in using experience from mature temporal cohorts to inform judgement on those less mature. However, without an appreciation of how these cohorts differ from each other, the provisioning exercise will be inherently flawed.

Data Used in Estimation Work

Overview

The data used to assess a variety of approaches to reserving during inflationary change are the same as those used in our previous analysis of inflation estimation methods *(Creedon, Bargate, Lenney, Schofield, & Stock, 2024)*. For the sake of completeness, however, the description of said pseudo-data and inflation scenarios is briefly repeated here.

For our purposes, it was desirable to create a dataset 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. More pertinently, the use of data considered to be "representative" of the real world, rather than drawn directly from it affords the advantage of control. In other words (and as alluded to above), inflation is the key variable of consideration in our work and we need not worry that changes in business mix, exposure, claim handling philosophy, etc. will be having an unduly distortive effect.

Claims derived using our tool 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 all publications.

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

Generation Method

As mentioned, VBA was used within Excel to generate:

- Fifty simulated years of individual claims type claims per trial; with
- A random number of claims in each accident (equivalently underwriting) year; and
- Each claim having a 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 may be found in the appendix.

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

The data generation tool was used to construct pseudo-real datasets used to gauge the efficacy of inflation estimation methods under various scenarios. It was desired to keep the majority of the parameters and assumptions 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.

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

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', rather than for any other particular reason. 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.

The random seed was set to 123,456.

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 so as to, again, better approximate 'real-world' data. Claims' development, for instance, does not consistently follow a 'nice' pattern in real life.

Claims were assumed to follow the development patten shown below. It was adjusted to be fully-developed after ten years, for ease of analysis.

Table 1: Large Claim Development Pattern

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

It should be reinforced that the data used for both the inflation estimation *(Creedon, Bargate, Lenney, Schofield, & Stock, 2024)* study and this exercise 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; etc.

Scenarios of Analysis

Overview

Our goal is, on a practical basis, to evaluate the various methods which may be used to accurately estimate claim provisions during periods of inflation volatility. In order to do so, we thought 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.

As described above, our data generation tool allows for (additive) inflation loads in economic, social and other/shock categories; as well as a trend of increasing/decreasing frequency. These inputs were used in setting the various scenarios and this section of the report briefly describes each of the scenarios under which inflation estimation methods were tested.

Scenario Descriptions

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.

Table 2: Total inflation values assumed for Scenario A

Calendar Year	2014	2015	2016	2017	2018	2019	2020	2021	2022	2023
Input	4.0%	4.0%	4.0%	4.0%	4.0%	4.0%	4.0%	4.0%	4.0%	4.0%

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 to 10 (inclusive). Thus, total inflation of claim amounts is 4% for years 1 to 4 and 6% for years 5 to 10, as well as future years thereafter (beyond the latest position of the triangle).

This scenario is designed to represent a trend of social inflation which is not initially present, but which manifests itself over time.

Table 3: Total inflation values assumed for Scenario B

Calendar Year	2014	2015	2016	2017	2018	2019	2020	2021	2022	2023
Input	4.0%	4.0%	4.0%	4.0%	6.0%	6.0%	6.0%	6.0%	6.0%	6.0%

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 to 4; 6% for years 5 to 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) fives years after the endpoint of the triangle.

Scenario C was designed to represent something of the state in which Casualty classes may have been argued to have found themselves shortly following the COVID-19 pandemic. In other words, inflation was reasonably stable historically; was said to be impacted by rising claims' costs due to social inflation in the recent past; and has (in theory) been subject to additional supernormal economic inflation in very recent years.

Table 4: Total inflation values assumed for Scenario C

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

Scenarios D1 and D2 – Non-Inflationary Frequency Trends

Scenarios D1 and D2 represent 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, claims' 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 claims' costs (i.e. severity) will be rising, but other external factors will be acting to increase or decrease the likelihood of claim 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 police spending are likely to result in an increase in crime rates and associated claim frequencies. All the while, general economic inflation will be acting to increase the cost of such claims when they do occur.

Scenario E – Unknown (to Analyst) Inputs

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, for the purposes of the original estimation exercise 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; and
- Second inflationary shock which emerge at 2% in year 9 and rises to 8% in year 10.

Total inflation parameters across all years are, thus, as follows:

Table 5: Total inflation values assumed for Scenario E

Calendar Year	2014	2015	2016	2017	2018	2019	2020	2021	2022	2023
Input	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.

Scenarios A through D2 formed the primary focus of the investigation. For each of these scenarios, reserving methods were compared with the generated ultimate/IBNR values to assess accuracy under the assumption that the "input" inflation was known to the user. However, in conducting the estimation study *(Creedon, Bargate, Lenney, Schofield, & Stock, 2024)* we further assessed how these methods would perform using the best estimate of the 'unknown' inflation parameters from the original estimation exercise.

High-level Methods for Incorporating Changing Inflation into Claims Reserve Estimates

Scene Setting

A wide variety of methods exist for the determination of best estimate claims reserves, with the topic having received highly extensive coverage in actuarial academic literature. With this being said, a handful of reasonably simplistic methods are pervasively employed in the setting of the overwhelming volume of reserves in the UK and London Market, namely:

A. The chain ladder (CL) method;
B. A priori (i.e. “initial expected”) methods; and
C. The Bornhuetter-Ferguson (BF) method - being, in essence a credibility-weighted blend of the former two.

It is expected that the reader will be reasonably familiar with these methods, though detailed descriptions may be found in many sources, such as the Institute and Faculty of Actuaries’ (IFoA) Claim Reserving Manual *(Faculty and Institute of Actuaries, 1997)*. These approaches are typically used in concert during a reserving (or even pricing exercise), with the choice of method typically depending on the maturity of a temporal cohort (origin year).

Although familiarity with these quotidian approaches is liable to be universal; it is worth reminding readers of one of the chain ladder method’s salient features. Namely, that the chain ladder inherently factors inflation into the setting of ultimate claim positions under the implicit assumption that future inflation will mirror that historically observed in the experience data. As such the triumvirate of CL, BF and a priori reserving techniques can be reasonably expected to sufficiently factor inflation into the best estimate under the proviso of stable inflation.

In listing these three methods, we do not assert that no other approaches to estimating claims reserves are in use within the London Market (or more widely in the UK and other jurisdictions). However, it is equally not unreasonable to note that the use of these approaches in setting claim reserves is highly ubiquitous.

Bearing this point in mind, we now list four approaches identified by the Corporation of Lloyd’s - *(Lloyd's of London, 2022)* and *(Lloyd's of London, 2022)* – to incorporating inflation volatility in best estimate claims reserving; methods which can be incorporated alongside existing techniques with relative ease. Each method will receive a brief description in this section and are listed thusly:

- A priori estimate uplift;
- Inflation-adjusted chain ladder method (IACL);
- Cashflow uplift; and
- Explicit management loading.

In the CL, inflation is implicitly factored in and hence no explicit assumption of claims inflation is required, provided inflation is believed to be highly stable. However, as will be noted below, a feature of inflation-adjusted methods is that they necessitate not only an estimate of historical claims inflation, but also a judgement as to its forecast. Accordingly, the need to both accurately estimate inflation present in an historical claim cohort and the relationship of this inflation with external drivers attains a considerable degree of importance.

A Priori Estimate Uplift

As mentioned, the BF and a priori methods are commonly used to set best estimate claim ultimate positions for temporal cohorts considered to be "insufficiently" mature vis à vis use of the CL method. What constitutes sufficient maturity will be a matter for the practitioner's judgement. The a priori reserving method, in essence, is (ipso facto) the selection of an ultimate position without reference to the actual experience of that temporal cohort.

Typically, the a priori approach takes the form of claims metric scaled to the exposure for that cohort. More specifically, the most common application of the a priori approach is to apply an estimated loss ratio to the ultimate premium estimate for the cohort to determine an estimate of ultimate claims. Hence, the a priori method is often referred to as the initial expected loss ratio (IELR) method.

A variety of augural approaches may be taken to selecting the a priori estimate. Most common, however, will be the use of historical experience, (either directly from the cohort of interest or from a relevant benchmark dataset) suitably on-levelled, to inform the selection.

Although the on-levelling of historical experience to inform the a priori selection will incorporate inflation, it is crucial to note that the application of this on-levelling will be on the basis of origin year (e.g. accident or underwriting), with the inflation associated with a given origin year being a weighted average of the inflation in future settlement years for claims of that origin year. The weights themselves will be determined by the payment or settlement pattern of that cohort. Accordingly, the inflation assumption for a given origin year will, by definition, be a blend of observed inflation in historical calendar years and assumed future calendar year inflation. This point is discussed in more detail in our previous paper *(Creedon, Bargate, Lenney, Schofield, & Stock, 2024)*.

With this in mind, it is arguable that a priori estimates or initially expected loss ratios ought to be adjusted over time to reflect observed calendar-year inflation and revised views on *future* calendar year inflation. This approach is described by the Lloyd's *(Lloyd's of London, 2022)* as an "*explicit loading applied to IELRs based on analysis of how inflation is expected to impact ultimate loss ratios"*. The word "loading" here was employed in their 2022 guidance letter on the basis that supernormal inflation was being experienced at that time and revision of originally estimated IELR was likely to be significant.

Cashflow Uplift

Uplifting cashflows is a reasonably straightforward approach to quickly factoring a loading for an abrupt inflationary change into claim reserve estimates. As such, the approach was viewed favourably by Lloyd's in their 2022 responses to the sudden emergence of supernormal inflation, following decades of stability *(Lloyd's of London, 2022).*

The salient advantage of the cashflow method is its relative simplicity and its comfortable marriage with whatever existing approaches are used by an organisation to estimate the ultimate claim positions/provisions. The essence of the approach is to recognise the implicitness of an inflation assumption within the CL (and by some extension BF and a priori) method; express a judgement as to what this implicit inflation is; and apply an uplift according to the difference between this future inflation forecast the estimate of the implicit historical assumption.

For a given cohort, the practitioner first produces a best estimate reserve position and the associated future cashflow pattern of those reserves. The individual cashflow values are then uplifted by an inflation factor (compounded over time) set with reference to the perceived difference between implicit historical inflation and judged forecast inflation.

A necessary extension of this method will be to determine the level of difference between average implicit (in CL) historical inflation and future inflation at the individual origin year level and apply a differential uplift to the cashflow associated with each origin year. Equally, this method should take into account any changing views on future inflation. E.g. if we expect future inflation to revert to the historical average, this should be taken into account in the projected cashflow uplifts.

Management Uplift

As the COVID pandemic began to recede in concern and supernormal inflation began to garner significant media attention in late 2021 and early 2022, (re)insurers in the US and Europe perceived as pressing, a potential for provision paucity. Reserving processes may not universally be (or have been at that time) sufficiently flexible as to be able to adapt to a phenomenon not seriously experienced in many decades. As such, a rapid solution to the risk of supernormal inflation resulting in reserve inadequacy was often considered to be an explicit, high-level loading – either to be held within the best estimate position or as an additional loading above the best estimate. We refer to such explicit additional provisions as “management loadings”.

The phrase “management loading” is naturally ambiguous. Unlike the three previously mentioned approaches, there is no hard and fast rule for setting an additional inflationary loading to be applied to any claims’ provisions (either within or on top of the stated best estimate). However, it is reasonable to assume that the aim of any such provision should be to reflect, with reasonable accuracy, the impact of the emergence of supernormal inflation (or, indeed, inflation volatility). As with the general principle of setting provisions on a best estimate basis, too low puts claimants at risk and too high risks being uncompetitive.

Accordingly, for the purposes of method evaluation, the Working Party considered the management uplift method to be, in essence, a simplification of the cashflow uplift method, namely:

a) Calculate the best estimate provisions as normal, under the assumption of unchanged inflation;
b) Calculate or estimate the duration or mean time to payment of the total best estimate provision; and
c) Apply an uplift to the best estimate provision (a) equal to the difference between assumed forecast and historical inflation, compounded by the duration (b).

In essence, this method can be seen as a simplified variant of the cashflow uplift approach. Rather than uplifting each individual cashflow (including by origin year), we apply a single uplift to the overall provision. The simplicity of this approach means it can be employed with rapidity and with minimal assumptions.

Inflation-Adjusted Chain Ladder

As mentioned, the CL, BF and a priori are a triad of commonly and concertedly used reserving techniques. Incorporating inflation adjustments into a priori assumptions will assist in ensuring that immature cohorts (i.e. those to which the IELR method is applied) more accurately reflect changing inflation, as well as going some way to improve the accuracy of BF-utilising cohorts. However, a priori changes will have no impact on reserves set using the CL and only partial impact on those set using the BF (which blends CL and a priori).

Accordingly, if practitioners do not wish to disregard the CL method entirely, it will be necessary to employ additional means to reflect varying inflation (when it is believed to be in occurrence). The inflation-adjusted chain ladder method is a commonly cited approach here. This is by no means a novel technique, first appearing in 1980 (Richards, 1980) and being thoroughly described in the Institute of Actuaries’ reserving manual (Faculty and Institute of Actuaries, 1997), as well as featuring in various syllabi for actuarial qualifications. The technique was also noted by Lloyd’s in 2022 (Lloyd's of London, 2022) as a means to reflect supernormal inflation in loss reserve estimates.

Being so thoroughly documented in various literatures, we will not devote significant time to describing the technique here. In essence, the technique can be summarised (with contrasted to the basic CL) as follows:

a) Obtain the incremental form of the claim triangle of interest;
b) Inflate the incremental values of this triangle according to the as-at date of the data;
c) Rebase the inflated incremental triangle on a cumulative basis;
d) Perform the basic CL on the inflated cumulative triangle;
e) Obtain all incremental claim amounts (past and future);
f) Inflate these incremental values in accordance with the chosen calendar year inflation index, including assumed inflation in future calendar years; and
g) Calculate the adjusted IBNR equal to the sum of future, inflated incremental values.

The IACL has the advantage of being a reasonably simplistic modification to existing, widely-used CL methods. In addition, it can painlessly be incorporated into the BF (with inflation-adjusted a priori assumptions). Furthermore, it works directly with a calendar-year inflation index, rather than needing to map such an index to an origin year basis.

High-level Method Evaluation

In this section, we discuss the results of applying the inflation-incorporating reserving methods mentioned in the preceding section against each of the scenarios of analysis considered. Firstly, though, we give some additional information on the practical steps we employed in our empirical evaluation.

A Note on Our Application of the Methods Described Above

Considering again, briefly, the generated data used in this work, we note that for each scenario, fifty simulations were run. For each simulation within each scenario, the four provisioning approaches were applied.

A general chain-ladder/BF approach was applied across all projection methods, with the BF method used for origin years considered to be under 70% developed. Chain ladder link ratios were calculated under both ten and four-year volume-weighted horizons.

Specific notes on the application of the four methods are as follows:

A. *A priori uplifts:* This method applies a combination of CL and BF to the claim development data, where the BF prior assumptions are loaded, as described above. No adjustment is made to the standard CL approach. The BF method is selected for years where development is greater than 70%.
B. *Cashflow uplift*: This method again begins with the approach described in (A), but with the Cashflow Uplift method applied to the IBNR cashflow amounts.
C. *Management Uplift:* The application of this method is identical to that of (B), but with the Management Uplift method (rather than the slightly more involved Cashflow Uplift method) applied to the total IBNR provision.
D. *IACL*: Here, we follow the same method as in (A), above, but with the chain ladder ultimate positions produced via the IACL, as opposed to the standard CL approach.

A crucial point to mention with regards our application of the various methods is that methods B, C and D (per list above) were performed in tandem with the *a priori uplift*. The essential logic here is that any given origin year will be projected using a pure chain ladder method, pure IELR (a priori) method, or credibility-weighted blend of the two (BF method). Adjusting IELR for changing inflation will, therefore, only “correct” the element of any origin years which use them.

Accordingly, methods B and C (above) were applied only to the cashflows which arose from a “pure” chain ladder approach. The resulting cashflows were then aggregated to form inflation-adjusted ultimate estimates and blended with the inflation-adjusted *a priori* estimates to form final, inflation-adjusted ultimate positions. The approach to employing the IACL follows a similar logic.

For each scenario, a relatively large number of reserve estimates are produced (estimate per simulation per method per link ratio approach). In each of the sections discussing the results for each scenario, we accordingly show a scatter plot comparing the estimated total incurred but not reported (IBNR) provision per simulation versus the known, simulated IBNR. Each point on the scatter plot will represent the error for an individual simulation. Lines on the plot will also show the average over all fifty simulations and a range of one standard deviation on either side of this average (the inter-tertial range, approximately).

This plot will be shown only for the method which produced the lowest average error across all fifty simulations. In addition to the scatter plot of individual simulation results for the "best-performing" method, we show a summary table of the average error (over all simulations) for all four approaches.

Furthermore, we show the average error under a "do nothing" approach. By this we mean, the average level of reserve mismatch (redundancy/deficiency) if a user were to make no explicit allowance for changing inflation and, instead, continue to apply the chain-ladder and BF methods under the assumption of 4% annual inflation in perpetuity.

Although the techniques were applied on two link ratio bases, as mentioned, we will show results on the "volume 4" basis only. The rationale here is that if a practitioner is aware of a change to the inflationary environment, they will likely reduce their reliance on more mature cohorts of data. Results have been produced on both link ratio bases, as well as on the basis of ultimate (as opposed to IBNR) deficiency/redundancy.

Whilst the full range of results and outputs are not displayed in this paper, the Working Party are very happy to provide these to the interested reader upon request.

Results and Commentary by Scenario

Scenario A – Stable inflation across all years

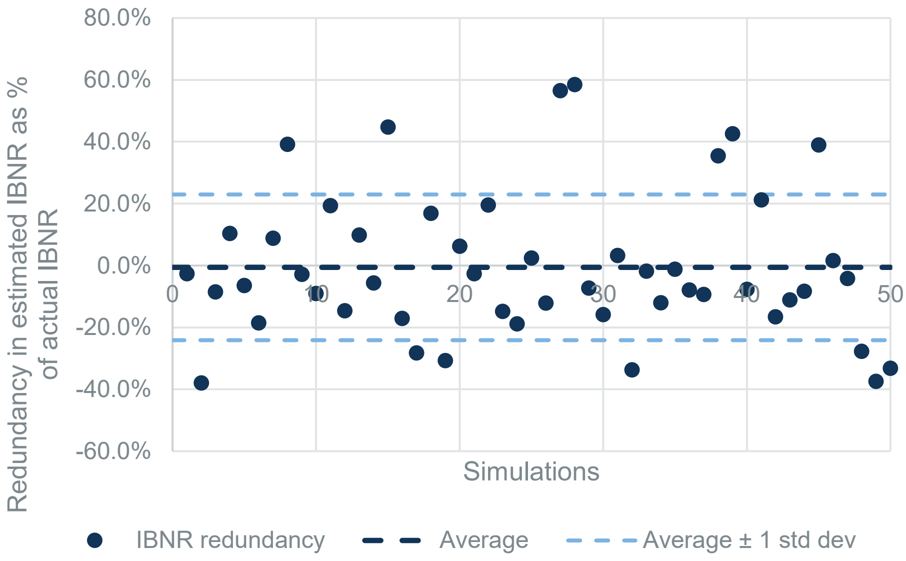

Figure 1: Scenario A - A priori method IBNR error

Table 6: Scenario A - Average IBNR error by method

	IBNR Deficiency (-) or Redundancy (+)
A Priori Uplift	-0.6%
IACL	-0.6%
Cashflow Uplift	-0.6%
Management Uplift	-0.6%
No Inflation Adjustment	-0.6%

Under the most basic scenario, all methods are, essentially, functionally equivalent and produce comparable results. Of, perhaps, some interest, however, is the observation that the level of error in any individual simulation can be quite significant. Although the average level of error across all fifty simulations is close to zero across all methods (see table above), process error at the individual simulation level can lead to notable reserve deficiency/redundancy.

Whilst it may be rather trite, this observation reminds us that a best estimate reserve position (typically) represents the mean of a probability distribution of possible outcomes and even with perfect knowledge of the form and parameters probability distribution, the final result will be unknowable and potentially subject to considerable process variability.

Scenario B – Emerging Social Inflation

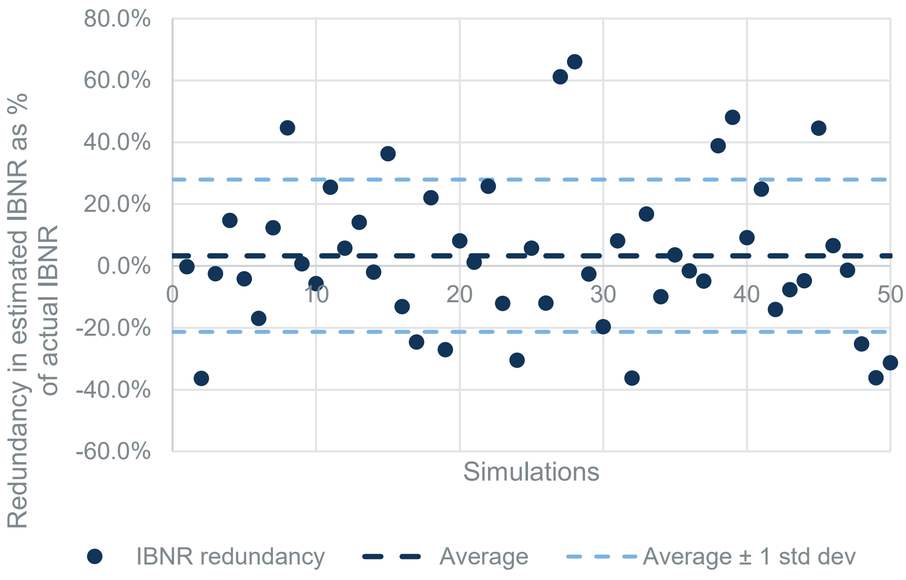

Figure 2: Scenario B - A priori method IBNR error

Table 7: Scenario B - Average IBNR error by method

	IBNR Deficiency (-) or Redundancy (+)
A Priori Uplift	3.3%
IACL	4.4%
Cashflow Uplift	3.3%
Management Uplift	3.3%
No Inflation Adjustment	-1.8%

In Scenario B, all methods, on average, estimate a greater level of IBNR than the true quantum. Scenario B is, again, a relatively simple case where inflation exhibits a persistent increase part-way through the period of analysis.

As one would expect, a failure to make any allowance for this inflation increase leads to a modest reserve deficiency. However, by utilising a relatively short averaging period in the CL method, we effectively take account of this change already.

All adjustment methods lead to relatively close estimates and small levels of redundancy. The IACL overestimates to a slightly greater degree.

Scenario C – Sudden Shock Inflation

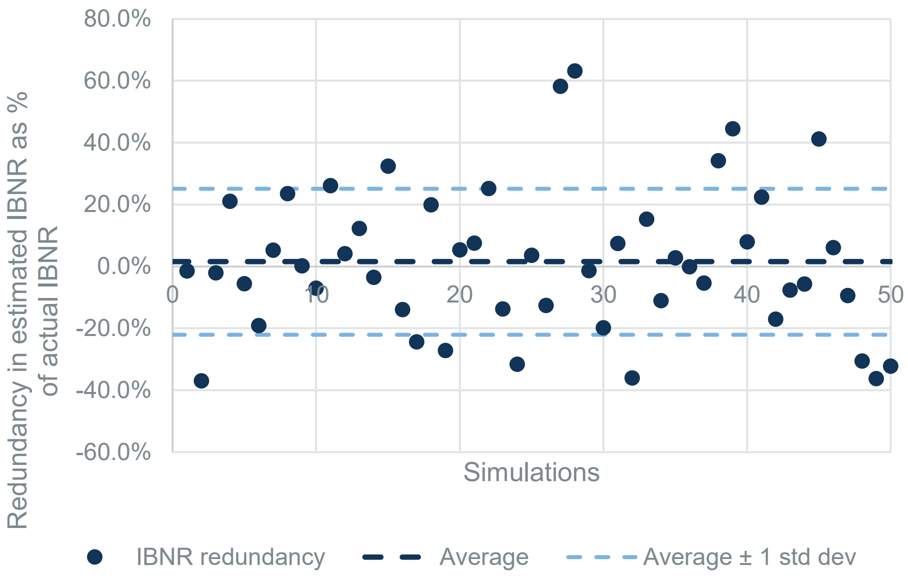

Figure 3: Scenario C – Management Uplift method IBNR error

Table 8: Scenario C - Average IBNR error by method

	IBNR Deficiency (-) or Redundancy (+)
A Priori Uplift	3.8%
IACL	11.0%
Cashflow Uplift	2.4%
Management Uplift	1.5%
No Inflation Adjustment	-5.6%

Scenario C represents a slightly more extreme view of inflation compared with Scenario B. Accordingly, the level of reserve deficiency when a practitioner makes no allowance for changing inflation is even more significant.

All adjustments methods are shown to overestimate the true level of ultimate. The cashflow and management uplift methods come closest to the correct result and are within under a percentage point of each other. These methods reflect the practitioner's assumed knowledge that the inflation spike has been reached and inflation is expected to trend downwards in future calendar years (i.e. after the youngest origin year for reserving). A pure, *a priori*, uplift, by contrast, overeggs to a slightly greater extent.

The IACL method, however, overestimates the true level of ultimate to a much more significant extent. We have found this overestimation to be reasonably ubiquitous across all scenarios, as will be shown over subsequent slides. We consider this to be a factor of how the IACL is poorer at compensating for the fact reversion of inflation to a lower level, following a spike. However, we consider this to merit further study.

Scenario D1– Decreasing Frequency Trend

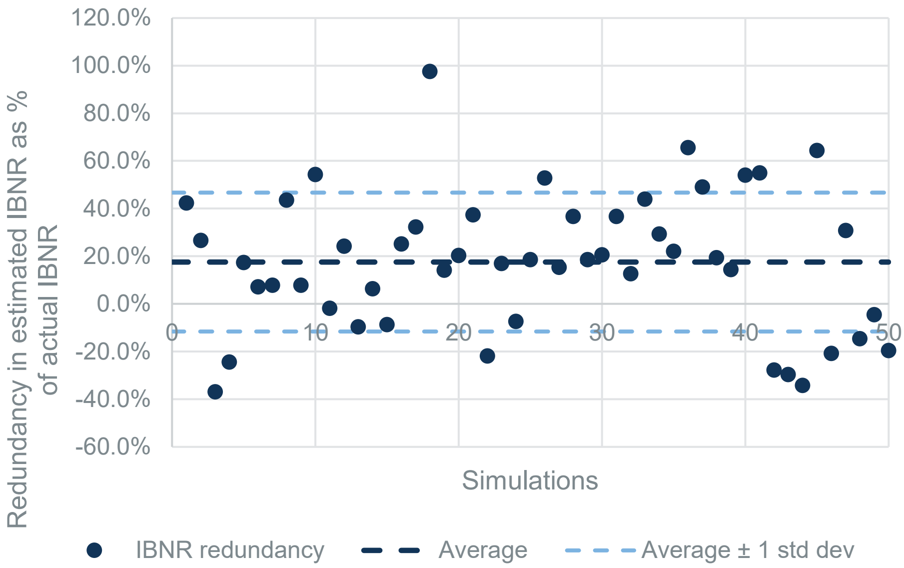

Figure 4: Scenario D1 - A priori method IBNR error

Table 9: Scenario D1 - Average IBNR error by method

	IBNR Deficiency (-) or Redundancy (+)
A Priori Uplift	22.0%
IACL	27.6%
Cashflow Uplift	19.1%
Management Uplift	17.6%
No Inflation Adjustment	12.5%

Unsurprisingly, all methods overestimate the true level of IBNR in the case of falling frequency. In this particular scenario, even a naive approach which makes no allowance for changing severity inflation leads to reserve redundancy; with this effect amplified for the methods which do attempt to adjust for rising inflation. As with Scenario C, the Management and Cashflow Uplift methods produce broadly equivalent results.

Scenario D2 – Increasing Frequency Trend

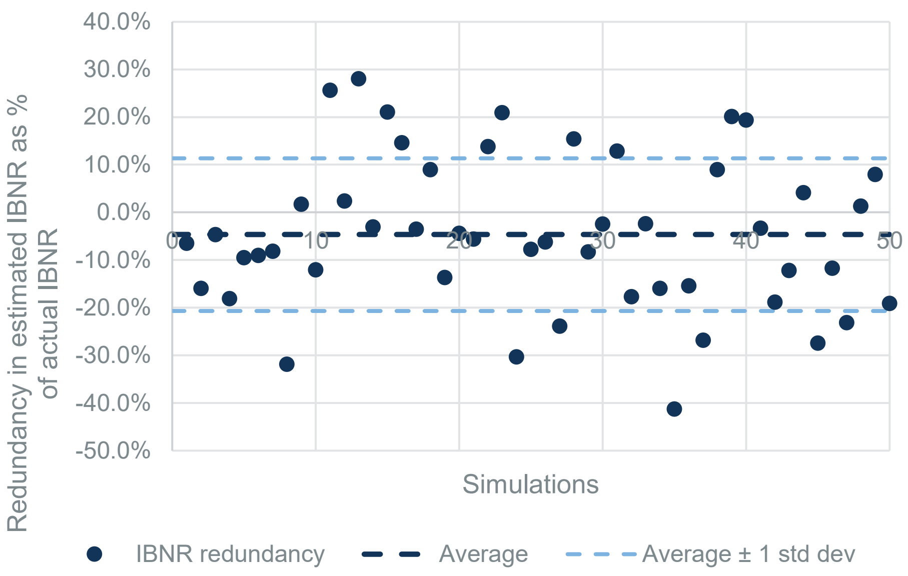

Figure 5: Scenario D2 - A priori method IBNR error

Table 10: Scenario D2 - Average IBNR error by method

	IBNR Deficiency (-) or Redundancy (+)
A Priori Uplift	-8.3%
IACL	-4.7%
Cashflow Uplift	-10.5%
Management Uplift	-11.5%
No Inflation Adjustment	-15.3%

In converse to Scenario D1, all approaches lead to a certain level of reserve deficiency in the case of increasing annual mean claim frequency; again with broad equivalence of the Management and Cashflow Uplift methods. Here we observe that the IACL adjustment leads to the closest match relative to true ultimate position. However, as alluded to previously, this is considered to be an artefact of a degree of prudence/bias in the method (or our application thereof) rather than an inherent ability to account for increasing claims' frequency.

Scenario E – Inflation Inputs Unknown During Estimation Exercise

For each of the previously-described scenarios – A through to D2 – analysis was performed on the basis of perfect knowledge of the inflation present in the historical data and that anticipated in the future. However, our original work on inflation estimation included a scenario in which the analyst performing the estimation exercise had no knowledge of the input parameters.

It was desirable to include a similar "unknown" scenario in our investigation on reserving approaches *(Creedon, Bargate, Lenney, Schofield, & Stock, 2024)*. As such, for Scenario E, inflation-adjusted reserving methods were conducted not with reference to the input inflation parameters, but rather the "most accurate" (as scored subsequently) estimate of inflation achieved by the analyst during that exercise.

Our objective here was to investigate how well the various reserving methods would perform in a scenario more akin to that of the real world. In other words, a scenario in which an analyst has far from perfect knowledge of the true level of claims inflation and is, instead, working with their best guess. Put another way, what is the compounding impact of parameter error?

As a reminder, below is a comparison of the inflation parameters used to generate the data for this scenario, versus those estimated by a user with access to the data, but no knowledge of input parameters. These estimates are a slight expansion of those presented in our previous paper *(Creedon, Bargate, Lenney, Schofield, & Stock, 2024)*, in that the "blind" user has used the Severity Trend method to determine an annual index (as opposed to period average).

Table 11: Scenario E input inflation parameters and blind parameter estimates

	2014	2015	2016	2017	2018	2019	2020	2021	2022	2023
Input Inflation	9.50%	10.50%	8.50%	6.50%	7.50%	7.50%	7.50%	7.50%	10.50%	16.50%
User Estimated Inflation	10.00%	6.00%	7.00%	7.00%	6.00%	6.00%	8.00%	8.00%	10.00%	10.00%

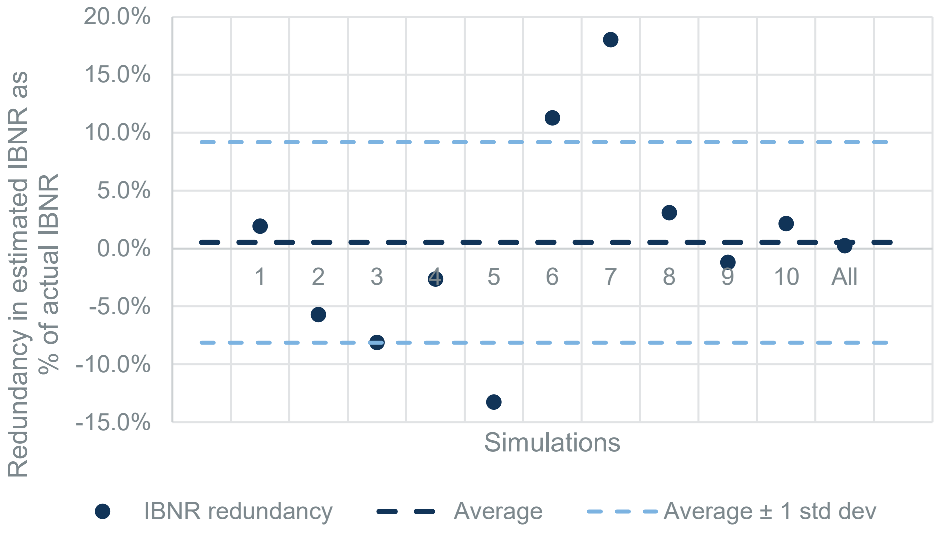

Figure 6: Scenario E2 - Cashflow uplift method IBNR error

Table 12: Scenario E - Average IBNR error by method

	IBNR Deficiency (-) or Redundancy (+)
A Priori Uplift	5.5%
IACL	12.1%
Cashflow Uplift	0.5%
Management Uplift	1.7%
No Inflation Adjustment	3.2%

As with previous results, we again observe close correspondence between the Management and Cashflow Uplift methods, in terms of average error level. What is, perhaps, surprising, however is that all methods (barring the IACL, as per previous) and even an unadjusted approach achieve a result within circa 5% of the true estimate.

Whilst at first glance, the closeness of the various methods (not to mention the unadjusted approach) may seem surprising, it is more understandable when we consider the specific features of the scenario. In particular, although the input inflation parameters used to generate the data are extreme; there is a significant offsetting impact of the 10% per annum negative frequency trend. It is considered that there would be merit in extending this particular exercise to a variety of additional scenarios, so as to better gauge the "parameter error" effect.

Overall Commentary

In general, we have observed that the various adjustment methods perform well with regards adjusting provisions to reflect a pure change in inflation. In particular, should there be an increase in claims' inflation vis à vis prevailing history; failing to make any allowance will lead to reserve deficiency, as opposed to applying specific adjustments. In general, though, the IACL was found to overstate the impact of rising inflation – either as an artefact of the method generally, or a nuance in our application thereof.

Crucially, however, adjustment methods which purely aim to deal with a change in inflation to claim quanta will naturally fail to account for any additional change in claim frequency. Although this result is to be expected, the impact is shown to be notable, even with modest frequency trend. Our subsequent section will consider lesser-used (in the London Market and UK) provisioning methods that make explicit allowance for such trends.

Additional Methods

Overview

The main focus of this paper thus far has been upon a relatively small number of approaches to address problems around inflation volatility/inflationary change in estimating loss provisions. The motivation for focussing upon these approaches has been both due to their relative ease of implementation and that they have been specifically highlighted in market guidance; as discussed previously. They also form relatively straightforward augmentations to the most ubiquitous reserving practices in the London Market. However, a variety of other approaches could feasibly be employed to better account for inflationary change in a reserving exercise.

Broadly speaking, we consider methods falling into two camps:

A. Stochastic or multivariate models with specific calendar year variables; and
B. Frequency/severity methods incorporating claim counts.

Multivariate stochastic reserving models have been the subject of much discussion in actuarial literature. Probabilistic trend family models also fall into this category. Owing to their relative complexity, they shall not be considered in explicit detail in this paper. However, the interested reader is encouraged to review Bohnert's 2015 paper, which explicitly considers these models from an inflation-risk perspective *(Alexander Bohnert, 2015).* Shi and Meyers have also penned a number of papers on the subject of multivariate stochastic reserving models, of which their 2014 paper, which explicitly considers calendar year effects (such as inflation) may be of interest *(Peng Shi, 2014)*.

By contrast, methods incorporating claim count data can be performed with relative ease. These approaches are comparatively uncommon in the London Market/US but see widespread use, particularly in Australia; as noted both in a 2024 presentation to the Casualty Actuarial Society, specifically considering inflationary challenges (Casualty Actuarial Society, 2024) and a discussion of these methods at the IFoA 2013 GI Pricing Seminar (Claughton, 2013). The latter presentation, in particular, offers interesting insights into the differences in approach and philosophy with regards provisioning in Australia vs. the UK/US.

In the remainder of this section, two such frequency-severity methods - commonly employed in Australia - will be discussed in further detail, namely:

1. The Payment per Claim Incurred (PPCI) method; and
2. The Payment per Claim Finalised (PPCF) method.

Of these, the PPCI method is perhaps best known in the UK. It is described in some detail in the IFoA manual on claim reserving *(Faculty and Institute of Actuaries, 1997)* where it is termed the "Bennett Taylor" method. It is also suggested by Lloyd's as a method for incorporating changing inflation into best estimate reserves *(Lloyd's of London, 2022)* under this name. The PPCF method can be considered an extension and refinement of the PPCI method. A concise description of these methods is provided below. They are discussed in much greater depth by their original proponent in *(Taylor, 2000)*. The most striking feature about these methods (from our perspective) is that, unlike the CL method, where inflation is implicitly allowed for, inflation allowances in these methods are always explicit.

It should be noted that the PPCI and PPCF methods require the availability of reliable claim count triangulations and, in the further case of the PPCF method, claim count finalisation/settlement triangulations. These data may not necessarily be readily available in a practitioner's organisation. Equally, the reliability of these data and their value in estimating loss provisions are not unquestioned. These points are considered by

Taylor and Xu in their paper which specifically considers the merits of incorporating claim count data in loss-reserving *(Greg Taylor, 2015).*

Description of Australian Methods

Payment per Claim Incurred

A salient feature of both the PPCI and PPCF method is the use of the incremental claim payment triangle, with each incremental value first inflated to the as-at date of the data. This step is performed as the methods do not natively build inflation into ultimate estimates, unlike the CL.

In broad terms, the PPCI method is performed as followed:

A. Estimate ultimate claim count by origin year (e.g. via CL, BF, etc.);
B. Calculate the ratio of inflated, incremental payment at each point in the paid claims' triangle to the ultimate claim count estimate for the applicable origin year;
C. Determine a selected (e.g. simple/weighted average similar to CL) ratio of those calculated in (B) per development period;
D. Complete the "rectangle" of the incremental, inflated paid triangle by multiplying the selections in (C) by the ultimate counts in (A); and
E. Inflate the future payments calculated in (D) to time of settlement.

Payment per Claim Finalised

The PPCF can be considered an extension of the PPCI method. The method goes further in that it explicitly considers settlement (finalisation) rates (also known as "operational time", per *(Reid, 1978)*), with the logic being that the bulk of a (non-structured) claim payment is paid at time of settlement. It requires a triangulation of settled claim counts.

Broadly, the PPCF method is performed thusly:

A. Estimate origin year claim count ultimate as per PPCI;
B. Calculate the ratio of cumulative claims settled in any cell of the settlement count triangle to ultimate count estimate by origin year;
C. Determine a selected settlement ratio of finalisations per development year (B);
D. Complete the "rectangle" of future finalisations by calculating the products of (A) and (C);
E. Calculate the ratio of inflated, incremental payment at each point in the paid claims' triangle to the incremental number of finalised claims in that same period (cell);
F. Determine the selected ratio of payment per claim finalised per (E) by development period;
G. Complete the "rectangle" of the incremental, inflated paid triangle by multiplying the selections in (F) by future incremental finalisations (D); and
H. Inflate the future payments calculated in (G) to time of settlement.

Assessment of Australian Methods

Quantitative Assessment

In their paper assessing the merits of claim count and finalisation count in determining loss reserves *(Greg Taylor, 2015)*, Taylor and Xu perform a thorough assessment of the PPCI and PPCF methods versus the CL against a variety of real-company data taken from the “Meyers-Shi” dataset *(G G Meyers, 2011)*. The assessment is highly rigorous and accompanied by detailed commentary around choice of dataset and model implementation.

The aim of that paper was to establish whether the use of claim count and claim settlement (finalisation) data via the PPCI and PPCF methods led to a superior fit, vis à vis the CL method. Although this proved to be the case in 80% of the datasets used, the focus of the paper was not explicitly on the impact of volatile inflation on loss reserve estimates. Accordingly, we performed a brief assessment of the application of the Cashflow Uplift method to the CL results derived in this paper, with the goal of ascertaining if this approach improved the fit of the basic CL method.

Note that the BF method could not be performed, owing to exposure data not forming part of the dataset. Equally, as we did not think to generate claim count triangulations in the pseudo-data used in our own exercises, the PPCI and PPCF methods cannot be performed on the data used to assess the four approaches described in the main body of this paper.

In order to perform the Cashflow Uplift exercise, it was first necessary to estimate the historical inflation levels present in the various company datasets provided. This was undertaken via a combination of the IACL and Separation methods approach described in our previous work *(Creedon, Bargate, Lenney, Schofield, & Stock, 2024)*. The results of this estimation exercise are shown in the chart below.

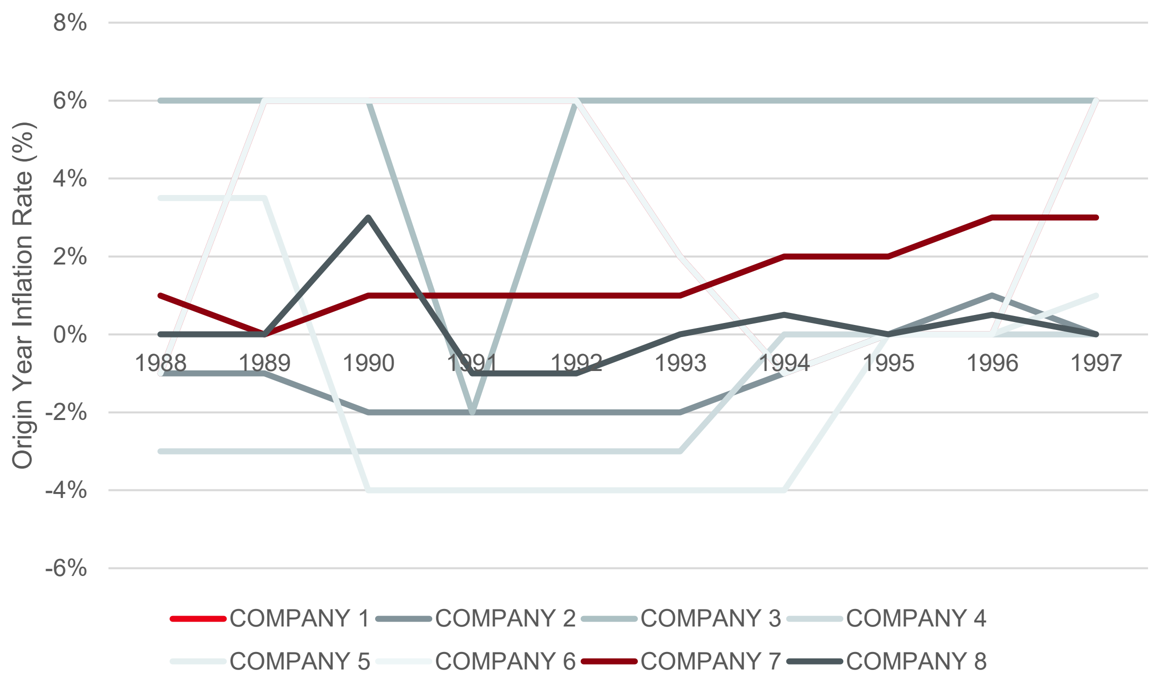

Figure 7: Estimates for Historical Claims Inflation by Company

As can be seen above, there is considerable variation in the empirically estimated claims inflation between each company. The estimates within an individual company are also often quite volatile. This volatility may well be an artefact of the fitting process, which was carried out on reasonably limited data.

Given these estimates of historical inflation, we next selected an uplift factor to be applied in the Cashflow Uplift method, based also upon a judgement of future inflation. The chart below then compares the level of total redundancy/deficiency (negative values) in the total (over all years) estimated reserves for the various company datasets assessed by Taylor and Xu, using their original CL, PPCI and PPCF methods against our Cashflow-Uplifted CL method.

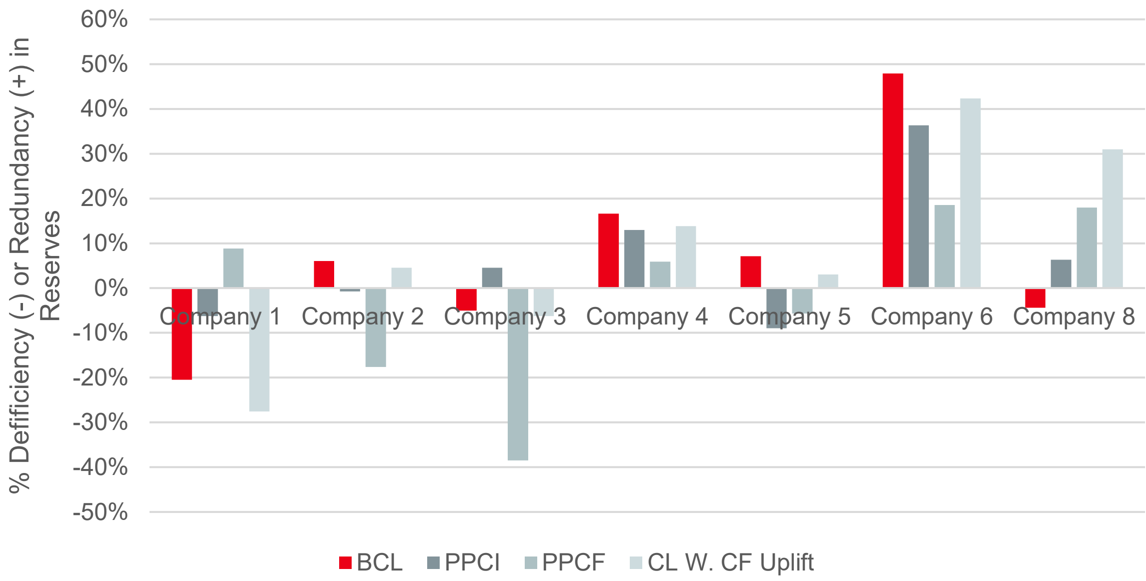

Figure 8: Reserve Redundancy/Deficiency (-) in Original vs. Cashflow Uplift Method by Company

With the exception of Company 8, the additional Cashflow Uplift is found to modestly improve the fit vis à vis the basic chain ladder (in terms of absolute error). However, barring Company 5, either the PPCI or PPCF outperforms the CL with Cashflow Uplift. The observed inflation for Company 5 is characterised by an uptick in the final origin years, which may perform well under the Cashflow Uplift method.

As mentioned above, our assessment here is reasonably brief and we do not intend it to be an exhaustive nor conclusive statement as to the relative merits of the claim count methods vs. CL with Cashflow Uplift; in particular due to exposure data being unavailable for this exercise. With this being said, varying inflation is a feature of each of the company datasets investigated and the Cashflow Uplift approach is shown to have improved the fit over the basic CL in almost all cases. Although methods incorporating claim count (either PPCI or PPCF) are still superior for the data above, it is questionable as to whether this is a result of incorporating the impact of variable inflation, or simply due to the additional predictive power achieved from incorporating claim count data.

We are enormously grateful to Dr. Taylor for providing us with the (neatly formatted) data and results used in his and Dr. Xu's paper.

Qualitative Assessment

The PPCI method, though infrequently used, is not unknown in the London Market, as mentioned above. The PPCF further refines the PPCI, through explicit consideration of settlement speed (operational time). Both methods are somewhat similar to the IACL in their approach to dealing with claims inflation. This is in the sense that they explicitly allow for known (estimated) historical and expected/judged future inflation on a calendar year basis.

Any additional predictive power associated with the PPCI/PPCF methods (vs. the four main adjustment methods discussed in this paper), therefore, may thus be postulated not to arise from a more accurate treatment of changing inflation, but rather from the additional information in the count/finalisation data. With this being said, the separate consideration of frequency/severity effects can readily be seen to be highly powerful. Failing to directly account for these, for instance, led to a high degree of inaccuracy in Scenarios D1 and D2 of our primary analysis.

Equally, it is a well-known feature that the presence of an excess or deductible in a (re)insurance coverage will distort the impact of claims inflation *(Simon Sheaf, 2005).* In the (reasonably common) case that loss amounts (i.e. claim severity) follow a Pareto distribution, however, it has been shown that severity in excess of this deductible remains unchanged in the presence of ground-up claims inflation *(Vytaras Brazauskas, 2009)*. In effect, the impact of the additional severity for claims "already" in excess of the deductible/excess point is offset by the additional frequency of claims previously below the deductible being "pushed" above it.

Although severity inflation will be flat in this case, there will be a noticeable trend in aggregate claims. However, an assessment of inflation levels present in excess (or non-proportional reinsurance) business which purely considered severity would miss this effect. Accordingly, the PPCI and PPCF methods may be a preferred approach by virtue of their separate consideration of frequency and severity.

Broadly speaking, there are two significant drawbacks to the to the PPCI and PPCF methods. Firstly, they hinge upon the availability of reliable count and finalisation triangulations, as mentioned previously.

Secondly, the methods are designed to be used in conjunctions with incremental paid claim triangulations. Accordingly, these methods neglect the useful information which is found in case estimates for open claims. This drawback could potentially be circumvented through the use of the incremental incurred (as opposed to paid) triangle. However, incremental claim payments very neatly lend themselves to the application of historical and future inflation on a calendar year basis. Incurred increments do not, as each incurred increment will represent a blend of payments and expected future payments.

Accordingly, the Projected Case Estimate (PCE) method also sees regular use in Australia. This method is briefly described in the appendix, with a more thorough overview provided by Taylor *(Taylor, 2000)*. The PCE approach behaves similarly to the CL applied to incurred claims, but differs in that it separately considers the case estimate (outstanding) component of the incurred position from the paid component. The method automatically produces projected payments, which are inflated to time of settlement. It does not, however, incorporate count data, unlike the PPCI and PPCF.

Ancillary Considerations

Temporal – What Inflation to Apply?

This issue was discussed in some detail in our previous paper *(Creedon, Bargate, Lenney, Schofield, & Stock, 2024)* but merits further consideration here, as it has significant implications for the application of loss reserving methods. In essence, claims provisions will invariably be calculated on an accident or underwriting year (collectively origin year) basis. However, inflation will manifest itself in the wider economy and be measured on a calendar year basis.

Claims occurring from a given origin year cohort will be paid over a variety of future calendar years. As such, the inflation associated with an origin year can be thought of as a weighted average of future calendar year inflation. These weights should be set equal to a payment, settlement or reporting pattern, as best judged by the practitioner.

We provide a brief illustration of this point below. 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:

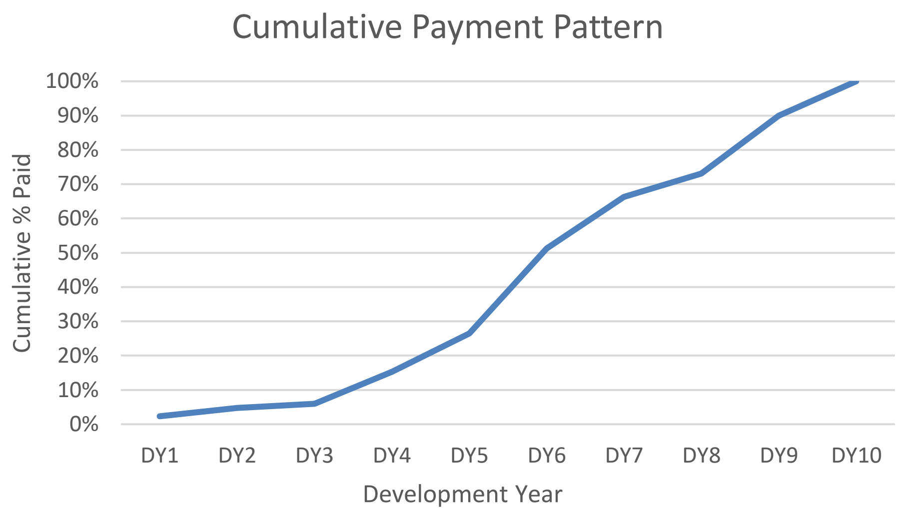

Figure 9: Example Payment Pattern (10Yr)

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. The chart below 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).

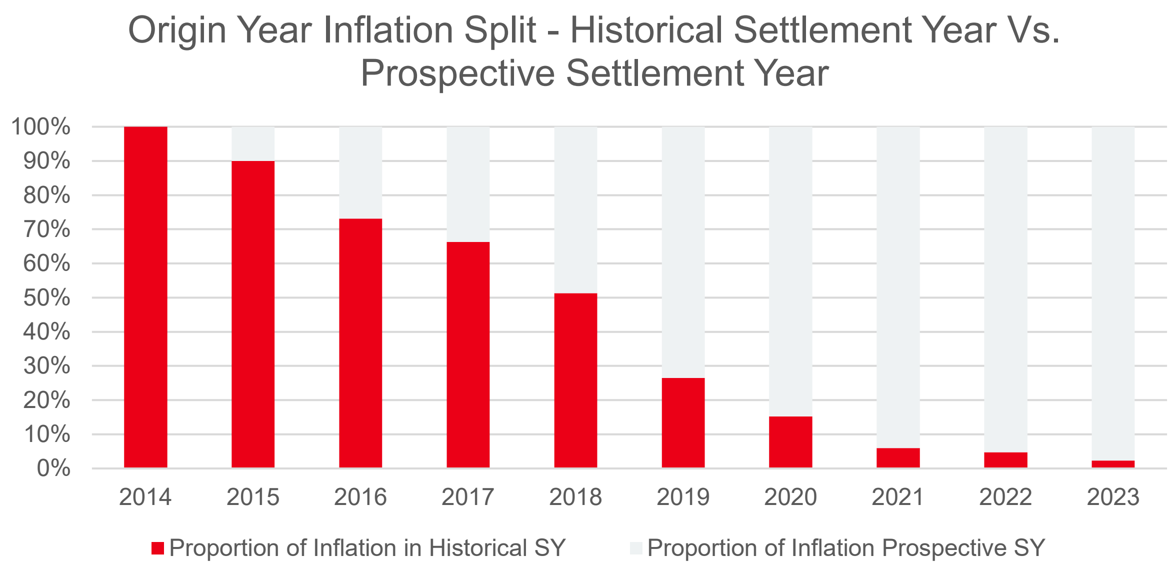

Figure 10: Origin Year Inflation Attributable to Past/Future Settlement Year Inflation

As can be seen, 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.

An implication of this thesis is that the inflationary change from one origin year to another represents the *expectation* of inflation in future calendar years. In uplifting a priori assumptions, as per suggested provisioning adjustment technique, we are revising this expectation to reflect the actual calendar year inflation emergence and revised views of future calendar year inflation.

Similarly, any application of the Cashflow or Management Uplift methods (to uplift CL or BF estimates) should take into account the following:

- Level of development and hence weighting applied to a priori assumptions in BF method (if used);
- Whether a priori assumptions are considered/have been amended to reflect future/emerged calendar year inflation; and
- Level of inflation assumed to be implicit in a given origin year CL ultimate vs. the assumed 'true' inflation for that origin year.

In essence, this points to a more complex Cashflow Uplift approach, whereby each origin year will need a separate uplift factor. The Management Uplift approach, though simpler, should also take this consideration into account in setting its uplift factor; even if performed heuristically.

As can be seen from the above, the popular (in London Market) trinity of CL, BF and a priori methods applied to incurred triangulations lead to a considerable degree of complexity when attempting to adjust for changing inflation, owing to this temporal point. Methods which avoid this complexity, such as the IACL (particularly applied to paid data) and Australian approaches which directly work with incremental payments strongly merit consideration in such a case.

Is Inflation Reflected in the Case Reserves Already?

Outstanding claim provisions or case estimates represent the expectations of future payments in respect of known claims. Case reserving practice and philosophy will vary heavily from firm to firm. For instance, it may be inadvisable for a (re)insurer to hold a reserve in respect of an advisory matter that has not yet been reported as a claim, for fear of adversely prejudicing final settlement. However, in theory, an individual case reserve should constitute the best estimate of liability for a claim reported to a (re)insurer.

Accordingly, an argument may be made that outstanding positions are already reflective of future inflation to time of claim settlement. In practice, the extent to which this is true may be questionable, as individual claim case estimates are likely to be updated only when there is a direct change in the status of that particular claim (e.g. receipt of a medical professional's assessment) and not necessarily in response to an external change in the inflationary environment. Equally, a move to revalue claims for changing inflation may be considered to introduce an unnecessary degree of uncertainty into these data.

In carrying out a provisioning exercise, practitioners should, however, be aware for the potential for this to double-count to occur. Careful consideration as to the make-up of reserves in terms of provision for claims yet to be reported (incurred but not yet reported or IBNYR) vs. provisions for underestimation of reserves on known claims (incurred but not enough reserved or IBNER) can assist here. Of greater importance is communication with the relevant claims function to ascertain the extent of any possible double-count with inflation-adjusted provisioning methods.

Inflation Dampening Policy Terms

Overview

This section pertains rather narrowly to Casualty Reinsurance. This line of business is characterised by the use of a number of policy terms which can act to dampen the effect of a rise in inflation on loss reserves. These terms include swing-rated premiums, sliding scale commissions and loss corridors on proportional business; as well as indexation clauses on non-proportional business. Sliding scale commissions and loss corridors are a feature of proportional business for lines other than Casualty; however they are particularly impactful (in respect of inflation) for longer-tailed lines.

A thorough description of policy features which may act to dampen unexpected inflation increases is found in the 2006 "Reinsurance Matters!" Working Party paper on Reinsurance Pricing *(Mark Flower, 2006).* The purpose of this section is to highlight the challenges such features pose in provisioning and so, only a very brief description is provided here. We split the policy features into two camps: indexation clauses and the others.

Swing-rated premiums, variable commissions and loss corridors all function in slightly different ways. However, the general principle associated with all of them is that as the performance of a reinsurance contract with such features deteriorates, said features will either act to increase the net (of commissions) premium payable to the reinsurer or reduce the losses paid by the reinsurer. The effect will be to (typically modestly) moderate the effect of an inflationary increase on the net (of commissions) loss ratio for that contract.

Indexation clauses function quite differently and feature in non-proportional reinsurance only. They appear commonly in Casualty (and Motor) reinsurance contracts in the UK and London Market, but feature less heavily elsewhere. Indexation clauses will either be "full" or "severe". In either case, an independent, third-party index (e.g. CPI) is specified at time of writing the contract.

In the case of a fully indexed contract, the original policy limit and attachment are indexed in-line with this benchmark from time of inception to time of underlying claim settlement, with the indexed limit and attachment then applied in determining reinsurance recoveries. Severe indexation clauses will further have a cumulative percentage assigned to them (e.g. 20%, 30%). The limit and attachment of the contract remain as-at inception until the cumulative value of the benchmark exceeds this percentage, at which the full level of indexation applies from then on.

Allowance in Reserving

An unexpected, significant uptick in the inflation applicable to a line of business (such as the supernormal economic inflation experienced in 2021 and 2022) will necessitate an upward adjustment to best estimate claim provisions, as has been noted. Equally, it may be surmised that the level of such an upward adjustment may be lessened by the existence of aforementioned policy terms and conditions in a portfolio. In the particular case of a severe indexation clause, the likelihood and speed of it being triggered may increase dramatically.

Although the above two points are reasonably without dispute, the question of *how* to allow for the benefit of policy terms and conditions when applying an uplift for inflationary increase poses challenges. Setting claim provisions is an exercise invariably carried out across a cohort of exposures. For a cohort of casualty reinsurance business, however, there will be a wide variety of terms and conditions in place. For a given cohort, the presence of such terms will vary from contract to contract, as will the precise formulation of such terms.

E.g., a carrier who writes a substantial UK Motor excess of loss account might reserve these contracts as a single cohort on the basis of their homogeneity. However, individual contracts within this cohort will likely contain a mixture of indexation clauses (full vs. severe at different thresholds). Equally, there will be differing order percentages, limits, and attachments across contracts. All of this will mean that a simple calculation of the effect of inflation-dampening features is unlikely to be possible.

Owing to the difficulty and uncertainty in calculating the benefit of such terms, it may be posited that giving credit for them in any best estimate provision would be unduly imprudent. However, there is still merit to attempting the calculation, so as to assist management in decision making and to better inform (and potentially reduce) any management loadings held in excess of the best estimate. Furthermore, the cynical observer might question the point of bothering to have a severe indexation clause to begin with, should a reserving function fail to give credit for its presence.

Two potential approaches could be adopted to achieve an approximation of the effect of these features: top-down or bottom-up. Both approaches will depend on possessing an accurate and systematic record of such terms.

A top-down approximation might proceed as follows:

- Determine the best estimate reserves for the cohort and associated adjustment for inflationary change or volatility;
- Perform an exercise to allocate the IBNR element of the claims ultimate to individual policy, split between unadjusted best estimate and inflationary load;
- Calculate the reduction to inflationary load at individual policy level based on this allocation; and
- Aggregate these reductions over all policies to determine the revised inflationary load for the cohort.

This may be a highly onerous series of calculations for a carrier to perform. Each policy term will act to dampen the inflationary adjustment in a different manner and so require its own formulation.

A bottom-up approach may instead take the following form:

- Estimate the proportion of IBNR applicable to policies with each relevant feature (fully indexed, swing-rated, etc.) across the entire cohort;
- Perform a sample calculation to estimate the reduction in inflationary load for a test policy for each feature; and

- Scale up in accordance with the estimate proportions.

This approach has the advantage of being reasonably transparent and simple to perform. The various judgements involved may also be flexed to provide a range. Naturally, though, the bottom-up approach is heuristic and not wholly rigorous. However, if the goal of the exercise is to provide a high-level view of the benefit of such policy features in reducing inflationary loads, this lack of rigour may be acceptable.

Closing Remarks

The purpose of this paper has been to assess techniques used to reflect changing or volatile inflation when estimating claim provisions. Given the preponderance of the Chain Ladder, Bornheutter-Ferguson and A Priori methods (most typically applied to incurred, rather than paid claim triangulations) in the London Market, our focus has been upon a number of approaches proposed by Lloyd's of London in 2022 to address this problem.

Our assessment of these techniques found that their application requires a considerable degree of care and judgement. In particular, all techniques require an estimate of both the inflation applicable historically to a given cohort and a forecast of future applicable inflation. Whilst it has been mentioned more than once in this paper, it merits reiterating that the inflation parameter associated with a given origin year (and which is used in on-levelling) is a blend of inflation over subsequent calendar years to time of claim finalisation.

Accordingly, the rate change and inflation assumptions used in on-levelling will necessitate revision, as our knowledge of historical and prospective calendar year inflation (vis à vis a given origin year) evolves over time. It was found that adjustment of a priori assumptions (in effect, IELR) went a considerable degree in producing estimates which matched the 'true' levels of provisions across our various scenarios analysed. However, the practice of continuously updating these prior assumptions for refined understanding of inflation may be a departure from existing practice for certain firms.

Equally, in the majority of our scenarios, we possessed perfect knowledge of historical and future inflation; permitting adjustments to provisioning techniques to be made with relative safety. The real-world practitioner will never have such knowledge and so the additional impact of parameter error (as per Scenario E) merits further research.

Alternatives to the four approaches proposed by Lloyd's were also considered in this paper. A variety of more complex methods exist which can directly incorporate inflation assumptions into the provision estimation process. These were not explored in any depth in this paper, but the interested reader is directed to consider them.

In addition, we also discussed two frequency-severity methods commonly used in Australia. The PPCI and PPCF techniques directly incorporate historical and future claim inflation parameters when setting provisions. Similar to the IACL, these techniques operate on a calendar year basis and, through their application on paid claim triangulations, elegantly circumvent the need to "allocate" calendar year inflation to origin year. Through separate consideration of frequency and severity, these methods can also more easily deal with disparate trends in these two components of claim provisions. However, although the PPCI and PPCF methods appear to be highly meritorious, they are not without drawbacks in their requirement for additional data and reliance on the paid triangulation.

With this being said, it was found that our methods to adjust provisions for varying (severity) inflation failed to capture the additional impact of frequency trend. To put it bluntly, it was akin to fitting a square peg to a round hole. Whilst widespread adoption of PPCI and PPCF techniques is not our unequivocal recommendation, it is clear that attempts at provisioning (particularly for long-tailed classes) will suffer if the potential for changes in frequency is not considered.

Inflation volatility or a significant change to the inflationary environment in which a class of business operates introduces additional challenges and uncertainty in setting best estimate claim provisions. A variety of reasonably efficacious techniques have been discussed to address this uncertainty. However, these techniques themselves require additional judgements and assumptions versus provision-setting during stable inflation.

The key additional assumptions needed are around the inflation parameters themselves and are largely common across all methods. The choice both of method and parameters can result in significant variations to the estimated provisions. More pertinently, inflation volatility, though continuing to be topical in 2025 remains but one source of uncertainty in determining provisions; as hinted at when we also begin to consider frequency trends.

If anything, our work has highlighted the extreme importance of an in-depth understanding of a given cohort of exposure or line of business. Without such an understanding, practitioners will consistently fail to provision accurately. Achieving such an understanding is not without challenges, as individual practitioners may work across multiple reserving lines, change roles, etc. Inflation aside, the greatest asset in improving accuracy of provisions may thus, rather, be a culture of open communication and professional challenge across the key stakeholders of claims, underwriting and the actuarial function.

Future Research

As mentioned at the beginning of this paper, a key question yet to be fully addressed by the Working Party is around the relationship between the observed inflation of claims for a given cohort (e.g. class of business) and externally tracked metrics or indices (e.g. components of price inflation). Establishing a relationship between cohort claims inflation and external metrics should permit a more rapid and more accurate revision of that claims inflation and hence better inform actuarial work around provisioning and pricing.

Although the relationship between claims inflation and external metrics will likely be highly bespoke to an individual line of business and geography, it is our hope that empirical investigations of these relationships may lead us to uncover some general principles on the subject also. We will look to continue this avenue of research in 2025.

Furthermore, we are keen to expand upon the general theme begun in this paper of the impact of changing inflation on day-to-day work. Whilst this paper addresses (in part) this query from the perspective of provisioning (arguably the starting point for much actuarial work in general insurance), there are key questions to address in pricing. We hope to consider such questions and the associated impact of *failing* to account for inflationary change in our future research.

Acknowledgements

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

- José Gómez Mena of Marcuson;
- Greg Taylor of the University of New South Wales;
- Dave Clark of Munich Re;
- Kate Bible of Aon, Australia;

Appendix 1 – Description of Claims Generation Tool

Overview of Section

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 in this appendix 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.

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.

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). 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 = 214013

    c = 2531011

    m = 2 ^ 32

    rd_num = (a * RdSeed + c)

    int_div = Int(rd_num / m)

    rand_num = rd_num - int_div * m

End Function
```

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 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.

Generation of Pseudo Claims

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

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.

Large Individual Claims

These are modelled in a similar manner to the attritional claims, with an intermediate step which 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 2 – Projected Case Estimate Method

Here is a brief overview on the application of the PCE method:

A. Obtain the triangulation of case reserve estimates and uninflated, cumulative (as distinct from PPCF and PPCI methods which use inflated incremental) triangulation of paid claims;
B. Calculate the ratios of change in cumulative payment (i.e. incremental payment) to the opening outstanding position, i.e. $PO_{i,j:j+1} = (P_{i,j+1} - P_{i,j})/OS_{i,j}$ for each applicable cell in the two triangles;
C. Determine a selected paid to outstanding ratio (PO) per development period increment from (B);
D. Determine a triangle of case estimate development (CED) factors as the ratio of the cumulative incurred point in a triangle cell less the outstanding in the previous development period to the cumulative paid in the previous development period $CED_{i,j:j+1} = (P_{i,j+1} + OS_{i,j+1} - OS_{i,j})/P_{i,j}$;
E. Determine a selected CED ratio per development period increment from (D);
F. Set the future incremental payment amount in each cell as the product of the outstanding case estimate in the previous development period with the selected paid to outstanding ratio from (D) i.e., $P_{i,j+1} = OS_{i,j} * PO_{i,j:j+1}$;
G. Set the future (projected) case estimate in any cell equal to the product of the case estimate for the previous development period with the selected case estimate development ratio (E) less the incremental paid for that period i.e., $OS_{i,j+1} = OS_{i,j} * CED_{i,j:j+1} - P_{i,j+1}$; and lastly
H. Inflate the future payments calculated in (F) to time of settlement.

The PCE method does not make use of claim count data and hence does not explicitly consider frequency trend, unlike the PPCI and PPCF methods. In addition, its use of case estimate development (CED) ratios engenders a degree of implicit inflation in the approach, similar to the basic CL method.

Bibliography

Alexander Bohnert, N. G. (2015). Assessing inflation risk in non-life insurance. *Insurance: Mathematics and Economics*.

Casualty Actuarial Society. (2024). CS-8 Double inflation (Micro + Social) - What Can We Do? Boston: Casualty Actuarial Society Reinsurance Seminar.

Claughton, A. (2013). Australian Reserving Methods. *General Insurance Pricing Seminar.* Institute and Faculty of Actuaries.

Creedon, C., Bargate, E., Lenney, S., Schofield, M., & Stock, R. (2024). *Claims Inflation Estimation: A Pracitcal Guide for Historical Data.* London: Institute and Faculty of Actuaries.

Faculty and Institute of Actuaries. (1997). *Claims Reserving Manual.*

G G Meyers, P. S. (2011). The retrospective testing of stochastic loss reserve models. *Casualty Actuarial Society E-Forum,*.

Greg Taylor, J. X. (2015). An empirical investigation of the value of finalisation count information to loss reserving. *University of New South Wales Business School*.

Lloyd's of London. (2022). Claims Inflation - What are Lloyd's Doing? Liverpool: Institute and Faculty of Actuaries.

Lloyd's of London. (2022). *Reserving Guidance - Allowing for Inflation.* London: Lloyd's of London.

Mark Flower, e. a. (2006). Reinsurance Pricing: Practical Issues and Considerations. *2006 GIRO "Reinsurance Matters!" Working Party*.

Peng Shi, S. B. (2014). A Bayesian Log-Normal Model for Multivariate Loss. *North American Actuarial Journal*.

PRA. (2023, January 10). *Bank of England.* Retrieved from https://www.bankofengland.co.uk/-/media/boe/files/prudential-regulation/letter/2023/insurance-supervision-2023-priorities.pdf

PRA. (2024, January 11). Retrieved from Bank of England: https://www.bankofengland.co.uk/-/media/boe/files/prudential-regulation/letter/2024/insurance-supervision-2024-priorities.pdf

Reid, D. (1978). Claim reserves in general insurance. *Journal of the Institute of Actuaries*, 211-296.

Richards, W. F. (1980). Evaluating the Impact of Inflation on Loss Reserves. *Casualty Actuarial Society*.

Simon Sheaf, W. F. (2005). Claims Inflation - Uses and Abuses. *Institute and Faculty of Actuaries - GIRO 2005.*

Taylor, G. (2000). *Loss Reserving: An Actuarial Perspective.* Boston: Kluwer Academic Publishers.

Vytaras Brazauskas, B. L. (2009). When Inflation Causes No Increase in Claim Amounts. *Hindawi Publishing Corporation, Journal of Probability and Statistics*.

Institute and Faculty of Actuaries

London

1-3 Staple Inn Hall · High Holborn · London · WC1V 7QJ

Tel: +44 (0) 20 7632 2100 · Fax: +44 (0) 20 7632 2111

Edinburgh

Level 2 ·Exchange Crescent · 7 Conference Square · Edinburgh ·EH3 8RA

Tel: +44 (0) 131 240 1300 · Fax +44 (0) 131 240 1311

Oxford

1st Floor · Park Central · 40/41 Park End Street · Oxford · OX1 1JD
Tel: +44 (0) 1865 268 200 · Fax: +44 (0) 1865 268 211

Beijing

Level 14 · China World Office · No.1 Jianguomenwai Avenue · Chaoyang District · Beijing, China 100004
Tel: + +86 (10) 6535 0248

Hong Kong

1803 Tower One · Lippo Centre · 89 Queensway · Hong Kong
Tel: +11 (0) 852 2147 9418

Singapore

5 Shenton Way · UIC Building · #10-01 · Singapore · 068808
Tel: +65 8778 1784

www.actuaries.org.uk