


SPLICE: a synthetic paid loss and incurred cost experience simulator

Benjamin Avanzi¹, Greg Taylor^{2*}  and Melantha Wang²

¹Centre for Actuarial Studies, Department of Economics, University of Melbourne, VIC 3010, Australia; and ²School of Risk and Actuarial Studies, UNSW Australia Business School, UNSW Sydney, NSW 2052, Australia

*Corresponding author. E-mail: gregory.taylor@unsw.edu.au

(Received 21 October 2021; revised 13 April 2022; accepted 14 April 2022; first published online 23 May 2022)

Abstract

In this paper, we first introduce a simulator of cases estimates of *incurred losses* called SPLICE (Synthetic Paid Loss and Incurred Cost Experience). In three modules, case estimates are simulated in continuous time, and a record is output for each individual claim. Revisions for the case estimates are also simulated as a sequence over the lifetime of the claim in a number of different situations. Furthermore, some dependencies in relation to case estimates of incurred losses are incorporated, particularly recognising certain properties of case estimates that are found in practice. For example, the magnitude of revisions depends on ultimate claim size, as does the distribution of the revisions over time. Some of these revisions occur in response to occurrence of claim payments, and so SPLICE requires input of simulated per-claim payment histories. The claim data can be summarised by accident and payment “periods” whose duration is an arbitrary choice (e.g. month, quarter, etc.) available to the user. SPLICE is built on an existing simulator of individual claim experience called SynthETIC (introduced in Avanzi *et al.* 2021b, *Insurance: Mathematics and Economics*, **100**, 296–308), which offers flexible modelling of occurrence, notification, as well as the timing and magnitude of individual partial *payments*. This is in contrast with the *incurred losses*, which constitute the additional contribution of SPLICE. The inclusion of incurred loss estimates provides a facility that almost no other simulators do. SPLICE is a fully documented R package that is publicly available and open source (on CRAN). SPLICE, combined with SynthETIC, provides 11 modules (occurrence, notification, etc.), any one or more of which may be re-designed according to the user’s requirements. It comes with a default version that is loosely calibrated to resemble a specific (but anonymous) Auto Bodily Injury portfolio, as well as data generation functionality that outputs alternative data sets under a range of hypothetical scenarios differing in complexity. The general structure is suitable for most lines of business, with some reparameterisation.

Keywords: Case estimates; Granular models; Incurred loss; Individual claims; Individual claim simulator; Loss reserving; Partial payments; Simulated losses; Superimposed inflation; SynthETIC; Synthetic losses

JEL codes: C52; C53; C63; G22

1. Introduction

1.1 Background and motivation

Machine learning methods are developing at increasing speeds in the actuarial literature. While the ultimate objective of these machine learning methods is application to real data, availability of synthetic data containing features commonly observed in real data is useful for at least two reasons: (i) such data sets, especially of granular nature and of large size, are in short supply in the

actuarial literature (see, e.g., Section 2.3 of Embrechts & Wüthrich, 2022), (ii) knowledge of the data generating process (impossible with real data) assists with the validation of the strengths and weaknesses of any new methodology.

Referring to scarcity of data (item (i) above), Embrechts & Wüthrich (2022), Section 2.3 mention two stochastic scenario generators: Gabrielli & Wüthrich (2018), and Avanzi *et al.* (2021b); Avanzi *et al.* (2021c). These simulators (and others, see for instance Section 5 of Avanzi *et al.*, 2021b, for a comprehensive review) produce synthetic claims *payments* experience. With the exception of the simulator developed by the “ASTIN Working Party on Individual Claim Development with Machine Learning” (Harej *et al.*, 2017), none of those simulators consider case estimates of incurred losses; see also Section 1.2 for further detail. These can be of significance for inference and prediction. An earlier reference on that point is Teugels & Sundt (2004), Volume 3, p. 1383 *et seq.*; see also Mack & Quarg (2004) on Munich Chain Ladder, and Taylor *et al.* (2008), who exemplify improved forecast performance when case estimates are taken into account.

The present paper describes an extension of SynthETIC, called SPLICE (Synthetic Paid Loss and Incurred Cost Experience), whose purpose is to close this gap in a manner consistent with the earlier paid claim experience, while leaving feature control on the hands of the user. In three modules, case estimates are simulated in continuous time, and a record is output for each individual claim. Revisions for the case estimates are also simulated as a sequence over the lifetime of the claim, in a number of different situations. Some of these revisions occur in response to occurrence of claim payments, and so SPLICE requires input of simulated per-claim payment histories. Furthermore, some dependencies in relation to case estimates of incurred losses are incorporated, particularly recognising certain properties of case estimates that are found in practice; a full list of our modelling overarching principles is provided in Sections 2.1 and 4.8. For example, the magnitude of revisions depends on ultimate claim size, as does the distribution of the revisions over time. The claim data can be summarised by accident and payment “periods” whose duration is an arbitrary choice (e.g. month, quarter, etc.) available to the user.

Although the three additional incurred loss modules in SPLICE could theoretically be used with any claims occurrence, notification and payment base (with care), we chose to make SynthETIC a required package, and use its existing structure without modifications. SynthETIC simulates paid claim experience of individual claims at a transactional level (key dates associated with a claim – e.g. settlement date – and claim payments). It offers an extremely flexible and complex base already: its modules represent specific features of the paid claim experience (e.g. claim sizes), and their plug-in nature hands control of these features to the user. Furthermore, when considering the output of SynthETIC along with those of SPLICE, the transactional simulation output now comprises key dates, and both claim payments and revisions of estimated incurred losses, and all this in a coherent and flexible format. This justifies the name of the package introduced in this paper: Synthetic Paid Loss and Incurred Cost Experience (SPLICE).

Finally, we refer to reason (ii) above, which mentioned the usefulness of synthetic data for model development and validation; see also Section 1.2.2. The *Annals of Actuarial Science* requires that submissions to its Actuarial Software stream demonstrate the analysis workflow using both synthetic and real data, highlighting the importance of synthetic data in this regard. When using SynthETIC and SPLICE together, the user has full control of the mechanics of the evolution of an individual claim. In particular, the user can decide the level of dependencies to include between different claim variates and test the effectiveness of any proposed new model in detecting such interactions. Indeed, by testing the proposed model against data across a spectrum of complexity, the user may derive new insights into its strengths and weaknesses, which is one main advantage of using synthetic data over real data. This is developed in Section 5.2.

1.2 Relation to prior literature

1.2.1 Claim simulation literature

Avanzi *et al.* (2021b) discussed a few predecessor simulators to some detail: Harej *et al.* (2017), Gabrielli & Wüthrich (2018), Coté *et al.* (2020), CAS Loss Simulation Model Working Party (2007), CAS Loss Simulation Model Working Party (2011), Bear *et al.* (2020). Of these, only the last three have been shown to simulate case estimates. All three generate samples of both paid and incurred amounts.

CAS Loss Simulation Model Working Party (2007) was superseded by the more advanced CAS simulator (CAS Loss Simulation Model Working Party, 2011). Here, case estimates are generated over the lifetime of each claim in continuous time.

Dates of incurred loss revisions are assigned randomly and independently over the claim lifetime according to a specified distribution. In most cases, a revision at (continuous) development time j is generated in the form

$$\text{Estimate of incurred loss} = \text{ultimate claim size} \times \text{adequacy factor},$$

where the adequacy factor is drawn from a log normal distribution. The parameters of this distribution may be defined by the user at fractions 0%, 40%, 70% and 90% of the claim lifetime. For intermediate fractions, the simulator interpolates the log normal mean. Drawings of distinct revisions appear to be stochastically independent.

This last feature ensures that the magnitudes of the revisions over the claim lifetime can be controlled. However, these revisions are independent of the size of the claim itself. SPLICE incorporates dependency in this respect. It also distinguishes between major and minor revisions, with further dependencies between the magnitudes of multiple revisions in respect of the same claim.

CAS Loss Simulation Model Working Party (2011) allows for inclusion of inflation according to accident period and, optionally, calendar period. In the latter case, the rates of accident and calendar period inflation are related. SPLICE allows the specification of arbitrary rates of inflation of either type. Calendar period inflation is specified as base inflation plus a superimposed inflation component. The rate of superimposed inflation may vary from claim to claim according to claim attributes such as ultimate size.

The more recent claim simulator sponsored by the Casualty Actuarial Society (Bear *et al.*, 2020) is structured differently from CAS Loss Simulation Model Working Party (2011). It appears to be concerned more with the simulation of the ultimate individual costs of a given portfolio of claims than with simulation of the detailed development of each claim.

It contains four options for the simulation of ultimate incurred loss, but only one of these generates a series of case estimates over the life of the claim. The other three simulate just ultimate incurred cost from the current claim status.

The one option that does simulate the development of incurred cost over claim lifetime does so by means of year-to-year development factors that are sampled from distributions defined by the user. These distributions differ from one development year to another, but the sampled development factors are stochastically independent, and there is no apparent provision for them to depend on the existing claim status (e.g. total paid to date). SPLICE remedies this.

In summary, SPLICE provides the following enhancements:

- Major and minor revisions are differentiated.
- The frequencies and magnitudes of these revisions can be made to depend on claim attributes such as ultimate cost.
- Dependencies are introduced between the magnitudes of revisions in respect of a single claim.
- The forms of inflation included are very general and flexible.
- Distributions of frequency and severity of revisions can be specified in a flexible way (a quality inherited from SynthETIC), that is, beyond lognormal. The complexity allowed here flows on to the incurred losses through their dependence on payment amounts and ultimate cost.

A brief mention of the rather different simulator of Côté *et al.* (2020) is appropriate here. This deals with a given set of data points of unspecified form, which might therefore, in principle, comprise time series of payments and case estimates for each claim.

The purpose of the suggested algorithm is to generate a synthetic data set with the same stochastic properties as the original. Generative adversarial networks are used to infer an underlying distribution of the data points, and then a new sample is drawn from this distribution.

The process is exemplified in Côté *et al.* (2020) using a well-known portfolio of French motor third-party liability policies (from CASdatasets, Dutang & Charpentier, 2019). The individual claim portfolio is re-sampled with respect to various claim attributes (car age, driver age, etc.) and claim count. However, no example demonstrating a resampling of paid claims and incurred claim costs is yet available.

1.2.2 Granular model literature

The development of so-called “granular (or micro-) models” requires availability of data at a certain level of detail. As such, such development is likely to benefit from the availability of simulator such as SPLICE, which generates this detail. De Felice & Moriconi (2019) provides a summary of the literature of these models, which are also discussed by Taylor (2019). However, the authors are unaware of any contributions to the granular model literature that consider the evolution of case estimates.

1.3 Package installation

SPLICE is released as an open-source R package on the Comprehensive R Archive Network (CRAN) at <https://CRAN.R-project.org/package=SPLICE> (Avanzi *et al.*, 2021a). In combination with SynthETIC, an existing open-source simulator of paid losses of individual claims available on CRAN (Avanzi *et al.*, 2021c), SPLICE simulates sequentially each of the 11 modules as outlined in Section 4, which provide the full functionality for the generation of synthetic paid loss and incurred cost experience (see Section 4). Its modular structure is discussed in Section 4.1.

SPLICE offers a collection of simulation functions for the incurred estimates. The default parameters for the simulation of revisions of incurred loss estimates are detailed in Appendix A, but they can be easily modified (or unplugged and replaced, if needed) by users to match their own experience with case estimates.

Users can choose to output their simulated claim histories in the form of a chain-ladder square of incurred losses by occurrence and development periods, or an individual transactional data set. In the latter, a transaction can be any of claim notification, settlement, a payment or a case estimate revision. A test transactional data set generated under the current specification is also available as part of the package (an example data excerpt is included in Appendix C). A full demonstration of the features offered by SPLICE can be accessed by running vignette("SPLICE-demo", package = "SPLICE") in the R console after the installation of the package.

Users can install the latest version of SPLICE from the CRAN repository via

```
> install.packages("SPLICE")
```

A development version of the program is also available on <https://github.com/agi-lab/SPLICE>. The GitHub repository contains, in addition to the package code, a chain-ladder analysis of the test data set discussed in Section 5.1, in an Excel spreadsheet, as well as some example datasets of various levels of complexity described in Section 5.2, available for download in .csv format.

1.4 Structure of the paper

The claim process is defined by the original eight paid loss modules (from SynthETIC) and an additional three incurred loss modules: major revisions, minor revisions and consolidation of

revisions, with the option to include inflation; see Section 4. SPLICE allows the user to specify any alternative form of distribution for the simulation of frequency, timing and magnitude of incurred loss revisions. The current default parameterisation has been set up to resemble the evolution of case estimates commonly found in practice; see also Section 4.8.

The general nature of case estimates is described in Section 2. After some notation in Section 3, SPLICE architecture is described in Section 4. Section 5 demonstrates the application of SPLICE, including an example implementation with the default parameterisation just mentioned, as well as an illustration of how alternative scenarios can be generated to achieve different levels of complexity. Section 6 contains some closing comments.

2. Case Estimates

2.1 General description

Most insurers assign **case estimates** to individual claims. A case estimate is here defined to mean an estimate of the ultimate cost of a claim, arrived at subjectively by means of expert knowledge. Case estimates are sometimes referred to as **manual estimates** or **physical estimates**. The experts who formulate them are usually known as **case estimators** or **loss adjusters**.

It is assumed that each claim carries, at each point in its lifetime, a case estimate of its ultimate incurred cost, and that the case estimators will vary these over time as additional information comes to hand.

It is assumed that revisions are either **major** or **minor**. Major revisions occur in response to material new evidence. For example, a claimant suffering head injury may be medically declared vegetative, in which case the perceived claim liability might increase substantially. Minor revisions occur as a result of more routine vagaries of a claim's progress. For example, unforeseen medical reports might be required.

Major revisions will be infrequent and usually of greater magnitude than minor. Moreover, major revisions represent a total change of perspective on ultimate claim cost, causing the case estimator to apply a revision factor to his estimate of that cost. Minor revisions, on the other hand, respond more to matters of detail, causing the case estimator to apply a revision factor to his estimate of outstanding payments.

The points below describe the development of a case estimate over the lifetime of a claim. Practice varies from one insurer to another, and the description given here may not fit all insurers. It would, however, describe a common practice.

The case estimate relates to the ultimate cost of the claim. Since the outstanding amount of the claim is equal to the difference between the ultimate cost and the paid losses to date, and since the latter is known at any point of the claim's lifetime, it follows that a case estimate of ultimate cost implies a case estimate of outstanding amount and vice versa.

The assumed features of incurred claims included in SPLICE are guided by the following overarching realistic principles:

Principle 1. *The insurer maintains case estimates of the incurred loss, and hence the outstanding loss, associated with each notified claim.*

Principle 2. *As long as there is no revision of the incurred loss, the estimate of outstanding loss is written down by each partial payment as it is made. This process is automated, and there is no intervention by the case estimator.*

Principle 3. *The case estimate of the incurred loss may undergo a number of revisions over the claim's lifetime.*

Principle 4. *These may occur at the time of a partial payment, or at any other time.*

Principle 5. *These revisions may be major (e.g. increase by a factor of 5) or minor (e.g. decrease by 5%).*

Principle 6. *By convention, each claim undergoes its first major revision at notification, when a case estimate is first established.*

Principle 7. *Major revisions other than this initial one are more likely for larger claims, and do not occur at all for the smallest claims.*

Principle 8. *They are relatively unlikely in the latter part of the claim's lifetime.*

Principle 9. *A claim may experience up to two major revisions in addition to the initial one, but the second, if it occurs at all, is likely to be smaller than the first.*

Principle 10. *Minor revisions tend to be upward in the early part of a claim's life, and downward in the latter part.*

Principle 11. *At settlement of the claim, the case estimate of ultimate cost will, by principle, be equal to actual amount paid, adjusted to the settlement date for base inflation. Correspondingly, the case estimate of outstanding claim cost will, by principle, be equal to zero.*

Although these are all listed as “principles”, it would also be fair to regard Principles 7–10 as design features. They are elevated to the status of “principles” here because they are, at least in the cases of Principles 7–9, commonly observed in practical claim portfolios. Principle 10 is somewhat different. Although it is encountered in many portfolios, alternatives are often encountered. The user should bear in mind that Principles 7–10 are, nonetheless, discretionary features of the default version of the simulator, and there is ample scope for their variation. For example, the user may choose to allow more than three major revisions over the course of the claim and make respective changes to the simulation of revision multipliers (which, by default, assumes a maximum of three major revisions).

2.2 Treatment of inflation

Here, **base inflation** is defined in Section 3.1 below as “normal” community inflation, such as price inflation or wage inflation, that would apply to claim sizes in the absence of extraordinary considerations. It is to be contrasted with **superimposed inflation** (“SI”), which represents the difference between the total rate of escalation of claim costs and base inflation. These principles are consistent with those of the original SynthETIC (Avanzi *et al.*, 2021c), but some additional assumptions and explanations are required with respect to the (new) incurred claims component.

Typically, case estimators are not expected to anticipate future base inflation. They may often be requested explicitly to exclude inflation beyond the valuation date. In such a system, each case estimate will represent ultimate claim cost in current-day values. This approach allows the insurance management to incorporate its own assumptions for future base inflation, which may depend on within-insurer consensus on economic conditions.

The estimators will usually be required to include full superimposed inflation up to the date of claim settlement.

Between case estimate revisions (see Principles 2–10), the estimated ultimate claim cost remains unchanged. Partial payments may occur, and the case estimate of outstanding claims will respond (see Principle 2), but the estimate of ultimate cost remains unchanged.

Revisions of incurred cost are the only points of intervention of the case estimators. At any such point, the estimator will adjust for base inflation to that point, i.e. an adjustment for the time elapsed since the immediately preceding revision. A “current-day value revision factor” will then be applied, representing the estimator’s change of opinion in ultimate claim size but with no allowance for any base inflation beyond the date of valuation.

3. Notation

3.1 Claim payments

The notation for claim payments was set out in Avanzi *et al.* (2021b). It is repeated here in Section 3.1 almost verbatim for convenience, as some will be required for the description of the incurred loss simulation. Its repetition will also provide a comprehensive view of the simulator in its entirety.

SPLICE works with **exact transaction times**, so time will be measured continuously. Calendar time $\bar{t} = 0$ denotes the first date on which there is exposure to occurrence of a claim. The time scale is arbitrary; a unit of time might be a quarter, a year, or any other selected period. The length of a period in years is specified by the user as a global parameter. The user needs to ensure that all input parameters are compatible with the chosen time unit.

For certain purposes (see Section 4), it will be useful to partition time into discrete periods. These are unit periods according to the chosen time scale. These periods will be of two types:

- **occurrence periods** (or accident periods), numbered $1, 2, \dots, I$, where occurrence period 1 corresponds to the calendar time interval $(0,1]$;
- **payment periods**, numbered $1, 2, \dots, 2I - 1$, representing the calendar periods in which individual payments are made, and including I past periods and a further $I - 1$ future ones.

An individual claim is settled by means of one or more separate payments, referred to here as partial payments. The claim will be regarded as settled immediately after the final **partial payment**. The delays between successive partial payments are referred to as **inter-partial delays**.

All payments are subject to inflation. They are initially simulated without allowance for inflation, and an inflation adjustment added subsequently. Any quantity described as “without allowance for inflation” is expressed in constant dollar values, specifically those of payment period 1. Inflation occurs in two types:

- (a) **Base inflation**: which represents, in some sense, “normal” community inflation (e.g. price inflation, wage inflation) that would apply to claim sizes in the absence of extraordinary considerations; and
- (b) **Superimposed inflation**: which represents the differential (positive or negative) between claim inflation and base inflation.

It is assumed that base inflation may be represented by a vector of quarterly inflation rates for both past and future calendar periods. The input inflation rates need to be expressed as quarterly effective rates irrespective of the length of calendar periods adopted. The inflation rates are used to construct an inflation index whose values are obtained:

- at quarterly points from calendar time 0, by compounding the quarterly rates; and
- at intra-quarterly points, by exponential interpolation between the quarter ends immediately prior and subsequent.

It is also assumed that SI occurs in two sub-types:

- (i) **Payment period SI**: which operates over payment periods; and
- (ii) **Occurrence period SI**: which operates over occurrence periods.

The following notation is used throughout:

- $\lfloor x \rfloor$ denotes the integral part of x
- $\lceil x \rceil$ = the ceiling function $\lceil x \rceil = \text{integral } n \text{ for } n - 1 < x \leq n$
- i = occurrence period $1, 2, \dots, I$
- \bar{t} = continuous calendar time with origin at the beginning of occurrence period 1
- $t = \lceil \bar{t} \rceil$ = payment period
- E_i = (annual effective) exposure in occurrence period i
- λ_i = expected claim frequency (per unit exposure) in occurrence period i
- $f(\bar{t})$ = base inflation index, representing the ratio of dollar values at calendar time \bar{t} to those at calendar time 0, constructed from the input base inflation rates
- $g_P(\bar{t}|s)$ = payment period SI index, representing the ratio of dollar values at calendar time \bar{t} to those at calendar time 0
- $g_O(i|s)$ = occurrence period SI index, representing the ratio of dollar values at occurrence period u to those at occurrence time 0
- n_i = number of claims occurring in occurrence period i
- r = identification number of claims occurring in occurrence period i ($r = 1, 2, \dots, N_i$)
- u_{ir} = occurrence time of claim r of occurrence period i (N.B. we have $i - 1 < u_{ir} < i$)
- s_{ir} = size of claim r of occurrence period i without allowance for inflation
- v_{ir} = delay from occurrence to notification of claim r of occurrence period i (N.B. the notification time is $u_{ir} + v_{ir}$)
- w_{ir} = delay from notification to settlement of claim r of occurrence period i (N.B. the settlement time is $u_{ir} + v_{ir} + w_{ir}$)
- m_{ir} = number of partial payments in respect of claim r of occurrence period i
- $s_{ir}^{(m)}$ = size of the m -th partial payment in respect of claim r of occurrence period i , $m = 1, 2, \dots, m_{ir}$
- $p_{ir}^{(m)} = s_{ir}^{(m)} / s_{ir}$ = proportion of claim amount s_{ir} paid in the m -th partial payment
- $d_{ir}^{(m)}$ = the inter-partial delay between from the epoch of the $(m - 1)$ -th to the m -th partial payment of claim r of occurrence period i , with the convention that $d_{ir}^{(0)} = 0$, corresponding to notification date (by convention, the 0-th “payment” is in fact the notification, without actual payment)
- $\bar{t}_{ir}^{(m)} = u_{ir} + v_{ir} + d_{ir}^{(1)} + \dots + d_{ir}^{(m)}$ = the epoch of the m -th partial payment

All of these quantities from n_i onward, but except r , are realisations of random variables. The random variables themselves are denoted in the same way but with the primary symbol in upper case. For example, S_{ir} denotes the random variable whose realisation is s_{ir} .

3.2 Incurred losses

The following additional notation, specific to case estimates, is introduced:

- τ = a generic variate denoting (continuous) time elapsed from claim notification
- $w_{ir}^{(m)} = d_{ir}^{(1)} + \dots + d_{ir}^{(m)}$ = delay from notification to epoch of m -th partial payment ($m = 1, 2, \dots, m_{ir}$) in the case of claim r of occurrence period i
- $w_{ir} = w_{ir}^{(m_{ir})}$ = delay from notification to settlement in the case of claim r of occurrence period i
- $w_{ir}^{(m_{ir}-1)}$ = delay from notification to epoch of the penultimate partial payment (i.e. the final major payment) in the case of claim r of occurrence period i
- $m_{ir\tau} = \text{largest integer } m \text{ for which } d_{ir}^{(1)} + \dots + d_{ir}^{(m)} \leq \tau$
- $c_{ir}(\tau) = \sum_{m=1}^{m_{ir\tau}} s_{ir}^{(m)}$ = cumulative claim payments up to and including delay τ from notification in respect of claim r of occurrence period i

- $y_{ir}(\tau)$ = case estimate of ultimate incurred loss at delay τ from notification in respect of claim r of occurrence period i
- $x_{ir}(\tau) = y_{ir}(\tau) - c_{ir}(\tau)$ = case estimate of outstanding claim payments at delay τ from notification in respect of claim r of occurrence period i
- k_{ir}^{Ma} = number of major revisions of incurred loss during the life of claim r of occurrence period i
- k_{ir}^{Mi} = number of minor revisions of incurred loss during the life of claim r of occurrence period i
- τ_{ir}^{Ma} = delay from notification to the epoch of l -th major revision of incurred loss during the life of claim r of occurrence period i
- τ_{ir}^{Mi} = delay from notification to the epoch of l -th minor revision of incurred loss during the life of claim r of occurrence period i
- g_{ir}^{Ma} = revision multiplier at the l -th major revision of incurred loss during the life of claim r of occurrence period i , causing $y_{ir}(\tau^-)$ to be replaced by $y_{ir}(\tau) = g_{ir}^{\text{Ma}} y_{ir}(\tau^-)$ at $\tau = \tau_{ir}^{\text{Ma}}$, where τ^- denotes $\lim_{\epsilon \downarrow 0} (\tau - \epsilon)$
- g_{ir}^{Mi} = revision multiplier at the l -th minor revision of outstanding claim payments during the life of claim r of occurrence period i , causing $x_{ir}(\tau^-)$ to be replaced by $x_{ir}(\tau) = g_{ir}^{\text{Mi}} x_{ir}(\tau^-)$ at $\tau = \tau_{ir}^{\text{Mi}}$

Finally, note:

- Revisions are assumed to occur at precisely the epoch τ , i.e. the revision has not occurred at τ^- .
- According to the explanation in Section 2.1, a major revision applies a factor to estimated incurred loss, whereas a minor revision applies a factor to estimated outstanding loss.

4. Architecture of the Claims Process

4.1 Modular structure

The claim process for claim r of occurrence period i is envisaged as consisting of the following modules:

Module 1: Claim occurrence date;

Module 2: Claim size without allowance for inflation;

Module 3: Claim notification date;

Module 4: Claim settlement date;

Module 5: Number of partial payments;

Module 6: Sizes of partial payments without allowance for inflation;

Module 7: Distribution of payments over time;

Module 8: Claim inflation;

Module 9: Major revisions of incurred losses (number of revisions, distribution of revisions over time and sizes of revisions);

Table 1. List structure of major revisions of incurred loss.

Name	Description
majRev_freq	Number of major revisions of incurred loss; see Section 4.2.1
majRev_time	Epochs of major revisions (time measured from claim notification); see Section 4.2.2
majRev_factor	Major revision multiplier of incurred loss ; see Section 4.2.3
majRev_atP	An indicator, 1 if the last major revision occurs at the time of the last major payment (i.e. second last payment), 0 otherwise; see Section 4.2.2

Module 10: Minor revisions of incurred losses (number of revisions, distribution of revisions over time and sizes of revisions);

Module 11: Development of case estimates, with the option to include inflation.

Modules 1–8 are present in the original version of SynthETIC (Avanzi *et al.*, 2021b; Avanzi *et al.*, 2021c) and were described there. Modules 9–11 are additional, and relate specifically to the simulation of case estimates in SPLICE. The present section details them.

Avanzi *et al.* (2021c) commented on the modular structure of SynthETIC. While the algebraic structure of SynthETIC has always been quite general, the authors noted that there could be cases where change of functional dependencies would be required. The modularity of SynthETIC has ensured that the user could unplug any one and replace with a version modified to his/her own purpose.

The modular structure of SynthETIC is retained in SPLICE. Each type of revision of incurred losses (major or minor) are simulated in three sub-modules: frequency of revisions, their distribution in time and sizes of revision factors; see Sections 4.2–4.3 and the subsections therein. The functional structure is designed to be general, and many users should be able to adopt it with changes to parameters but not algebraic structure. However, in cases in which some change of structure is necessary, the modules can be unplugged and replaced with ease.

The sequence of Module 1 to Module 11 must be preserved, because each module typically relies on the output of prior modules. Examples of such dependencies are provided in Section 5.1.3 (and in Section 4.3 of Avanzi *et al.*, 2021b).

4.2 Module 9: major revisions

This section introduces a suite of functions that work together to simulate, in sequential order, (1) number of major revisions of incurred loss (`claim_majRev_freq`), (2) distribution of major revisions over time (`claim_majRev_time`) and (3) factors of major revisions (`claim_majRev_size`), for each of the claims occurring in each of the occurrence periods.

In particular, `claim_majRev_freq()` sets up the structure of the output for major revisions: a nested list such that the j th component of the i th sub-list is a list of information on major revisions of the j th claim of occurrence period i . The “unit list” (i.e. the smallest, innermost sub-list that is unique to each individual claim) consists of the components in Table 1.

4.2.1 Number of major revisions of incurred loss (`claim_majRev_freq`)

The number of major revisions k_{ir}^{Ma} is the realisation of a random variable K_{ir}^{Ma} with df $F_{K|s}^{\text{Ma}}(k; s)$, specified on input as a function of k , and possibly dependent on claim size s .

The default version of $F_{K|s}^{\text{Ma}}(k; s)$ is set out in Appendix A. In addition to the default parametrisation, SPLICE supports a full range of user-specified alternative sampling distributions, with modifiable dependencies on other claim variates, as is the case for all modules that follow.

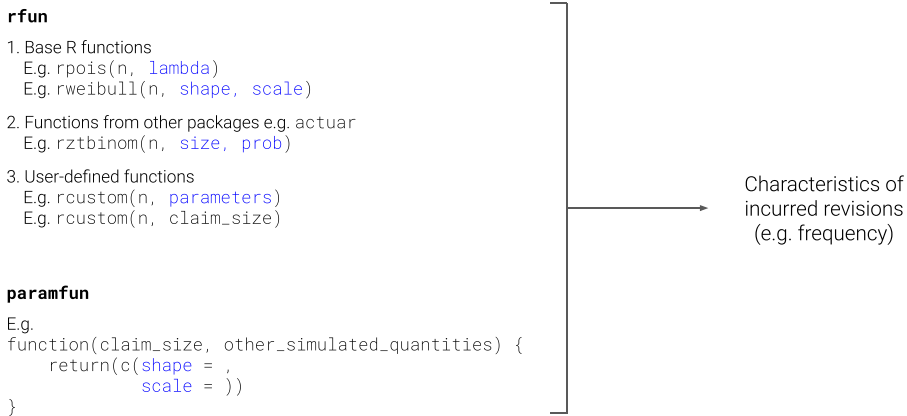


Figure 1. SPLICE implementation.

Users of the package can choose any suitable sampling distribution through the arguments `rfun` (a random sampling function) and `paramfun` (parameters for the random sampling function) to better serve their own testing purposes. As illustrated in Figure 1 below, `rfun` defines the functional form of the distribution and can be any of pre-defined distributions in base R, or more advanced ones from other packages such as `actuar` (Dutang *et al.*, 2008), or any proper user-defined function, while `paramfun` creates the link between the previously simulated quantities and the parameters of `rfun`. Here, “parameters” should be interpreted in a very general sense: for a specific parametric distribution – e.g., Weibull, this can simply be the shape and scale parameters of the distribution, defined as a function of already simulated quantities – e.g. `claim_size` (from `SynthETIC`). For a user-defined `rfun` taking any other arguments, `paramfun` should output the required `rfun` arguments as a function of the already simulated quantities.

The SPLICE vignette, which can be accessed via

```
R> vignette("SPLICE-demo", package = "SPLICE")
```

or online at <https://CRAN.R-project.org/package=SPLICE> (Avanzi *et al.*, 2021a), includes illustrative examples on using a zero-truncated Poisson distribution with both default and modified dependence structures (e.g. adding the dependence of K_{ir}^{Ma} on the number of partial payments m_{ir} of the claim in addition to its default dependence on claim size s_{ir}).

4.2.2 Distribution of major revisions over time (`claim_majRev_time`)

As noted in Principle 6, each claim experiences a major revision at notification. The following remarks relate to any subsequent major revisions. These occur only for claims with $m_{ir} \geq 4$.

If $l = k_{ir}^{\text{Ma}}$, that is, when considering the last major revision g_{ir}^{Ma} , the probability that this coincides with the penultimate claim payment (i.e. the final major payment) is $P(\tau_{ir}^{\text{Ma}} = w_{ir}^{(m_{ir}-1)} | s_{ir})$, possibly dependent on claim size. In this event (`majRev_atP == 1`), the $\tau_{ir}^{\text{Ma}}, l = 2, \dots, k_{ir}^{\text{Ma}} - 1$ are realisations of a random variable $T^{\text{Ma}1}$ with df $F_{T|w}^{\text{Ma}1}(\tau; w)$ specified on input as a function of τ .

In the event that the last major revision does not coincide with the penultimate claim payment (`majRev_atP == 0`), the $\tau_{ir}^{\text{Ma}}, l = 2, \dots, k_{ir}^{\text{Ma}}$ are realisations of a random variable $T^{\text{Ma}2}$ with df $F_{T|w}^{\text{Ma}2}(\tau; w)$ specified on input as a function of τ .

4.2.3 Factors of major revisions of incurred loss (`claim_majRev_size`)

As noted in Principle 6, each claim experiences a major revision at notification. The following remarks relate to any subsequent major revisions.

Table 2. List structure of minor revisions of incurred loss.

Name	Description
minRev_atP	A logical vector indicating whether there is a minor revision at each partial payment; see Section 4.3.1
minRev_freq_atP (minRev_freq_notatP)	Number of minor revisions that occur (or do not occur) simultaneously with a partial payment. minRev_freq_atP is numerically equal to the sum of minRev_atP; see Section 4.3.1
minRev_time_atP (minRev_time_notatP)	Epochs of minor revisions that occur (or do not occur) simultaneously with a partial payment (time measured from claim notification); see Section 4.3.2
minRev_factor_atP (minRev_factor_notatP)	Minor revision multiplier of outstanding claim payments for revisions at partial payments and at any other times, respectively; see Section 4.3.3

The factor by which estimated incurred loss is adjusted at a major revision g_{irl}^{Ma} is the realisation of a random variable G_{irl}^{Ma} with df $F_{G|h}^{Ma}(g; h)$, specified on input as a function of g , and possibly dependent on the history of major revisions h_l preceding the l -th.

The default version of $F_{G|h}^{Ma}(g; h)$ is set out in Appendix A. The SPLICE vignette illustrates how to deploy alternative sampling assumptions.

4.3 Module 10: minor revisions

For the simulation of minor revisions, we treat separately the case of revisions that occur simultaneously with a partial payment and the ones that do not (Principle 4 in Section 2.1).

Similar to the case of major revisions, the suite of functions under this heading run in sequential order to simulate (1) number of minor revisions of incurred loss (claim_minRev_freq), (2) distribution of minor revisions over time (claim_minRev_time) and (3) factors of minor revisions (claim_minRev_size), for each of the claims occurring in each of the occurrence periods.

Analogous to major revisions, claim_minRev_freq() sets up the structure of the output minor revisions: a nested list such that the j th component of the i th sub-list is a list of information on minor revisions of the j th claim of occurrence period i . The “unit list” consists of the components in Table 2.

4.3.1 Number of minor revisions of incurred loss (claim_minRev_freq)

The number of minor revisions k_{ir}^{Mi} consists of two components: $k_{ir}^{Mi} = k_{ir}^{Mi1} + k_{ir}^{Mi2}$, where those counted in k_{ir}^{Mi1} are simultaneous with a partial payment (possibly final payment), and those counted in k_{ir}^{Mi2} are not.

The variate $k_{ir}^{Mi1} = \sum_{l=1}^{M_{ir}} b_l$, where the b_l are realisations of independent Bernoulli variates B_l , each corresponding to the l -th partial payment, and having df $F_{B|l}(b)$.

The variate k_{ir}^{Mi2} is the realisation of a random variable K_{ir}^{Mi2} with df $F_{K|w}^{Mi2}(k; w)$, specified on input as a function of k , and possibly dependent on settlement delay w .

The default versions of $F_{B|l}(b)$ and $F_{K|w}^{Mi2}(k; w)$ are set out in Appendix A.

4.3.2 Distribution of minor revisions over time (claim_minRev_time)

If a minor revision occurs in conjunction with a partial payment, its epoch is equal to the epoch of the payment.

If a minor revision occurs at an epoch other than those of partial payments, then the τ_{irl}^{Mi} are realisations of a random variable T^{Mi} with df $F_{T|w}^{Mi}(\tau; w)$ specified on input as a function of τ .

Major and minor revisions cannot occur simultaneously. In the event that they are simulated to do so (which will only ever occur at the last major payment), the major revision takes

precedence, and the minor revision is discarded. This adjustment is made at the consolidation step (i.e. Module 11).

4.3.3 Factors of minor revisions of incurred loss (claim_minRev_size)

The factor by which case estimate is adjusted at a minor revision g_{irl}^{Mi} is the realisation of a random variable G_{irl}^{Mi} with df $F_{G|w,\tau}^{Mi}(g; w, \tau)$, specified on input as a function of g , and possibly dependent on the delay w from notification to settlement and the delay τ from notification to the subject minor revision.

The default version of $F_{G|w,\tau}^{Mi}(g; w, \tau)$ is set out in Appendix A. Again we refer to the SPLICE vignette and package documentation for illustrations of alternative parametrisations (Avanzi *et al.*, 2021a).

4.4 Module 11: Computation of case estimates (claim_history)

4.4.1 Without inflation

Initially, base inflation will be ignored. All case estimates will be computed in values corresponding to time $\bar{t} = 0$, i.e. the commencement of the first occurrence period.

For each claim, claim size (before base inflation) is simulated within the Payments section of SynthETIC. By Principle 11, the case estimate at settlement (again before base inflation) must coincide with it. Symbolically,

$$y_{ir}(w_{ir}) = s_{ir}. \tag{1}$$

In order to ensure this identity, it is necessary to simulate case estimates in reverse chronological time. One commences by setting $y_{ir}(w_{ir})$ in accordance with (1), then calculating $y_{ir}(w_{ir}^-)$. This will be equal to $y_{ir}(w_{ir})$ if no revision of incurred amount occurs at settlement. Otherwise, it will be calculated by means of (3) below.

Note that $y_{ir}(\tau) = y_{ir}(w_{ir}^-)$ for τ equal to the delay from notification to the epoch of the last revision (major or minor) strictly prior to settlement. From this $y_{ir}(\tau^-)$ is calculated, with allowance for the revision by either (2) or (3). Working in reverse order in this way, one calculates $y_{ir}(\tau)$ for all $0 \leq \tau \leq w_{ir}$. The value of $y_{ir}(0)$ arrived at is the initial case estimate (at notification) for the claim.

The relations used to calculate a pre-revision case estimate from post-revision case estimate at epoch τ are **initially** as follows:

$$y_{ir}(\tau^-) = \frac{y_{ir}(\tau)}{g_{irl}^{Ma}} \tag{2}$$

if the l -th major revision occurs at epoch τ ; or

$$y_{ir}(\tau^-) = c_{ir}(\tau^-) + \frac{y_{ir}(\tau) - c_{ir}(\tau^-)}{g_{irl}^{Mi}} \tag{3}$$

if the l -th minor revision occurs at epoch τ .

When a minor revision coincides with a partial payment, there is a need to define whether the revision of outstanding claims occurs first and is then followed by the payment, or vice versa. Note that (3) is equivalent to

$$y_{ir}(\tau) = c_{ir}(\tau^-) + g_{irl}^{Mi} x_{ir}(\tau^-), \tag{4}$$

which means that the revision occurs first.

As an illustrative example, Figure 2 visualises the development of two sample claims (without inflation). The grey paths in the plots are simulated by SynthETIC (Avanzi *et al.*, 2021b). They describe the full history of paid losses during the lifetime of a claim, which is taken as input by




Figure 2. Graphical representation of two sample claims from the example implementation.

SPLICE for the generation of the incurred histories (black paths). For reasons explained above, SPLICE works backward from the settlement of the claim, sets the case estimate at settlement equal to the claim size, and only updates the case estimates when a revision of incurred loss is projected to occur. The jumps in the black paths thus correspond to the points of major or minor revisions. We remark that some of the jumps may coincide with a partial payment (e.g. all three partial payments for claim #2 result in a simultaneous minor revision; left panel in Figure 2). The computation of case estimates in such cases is governed by equation (4), and detailed in Appendix C. For details of the two claim records, we refer to the data excerpt in Appendix C.

It is stated above that (2)–(4) apply only “initially”, because adjustments may be required to deal with constrained cases. For example, a large value of g_{irl}^{Ma} in (2) could force $y_{ir}(\tau^-)$ below $c_{ir}(\tau^-)$, in which case $x_{ir}(\tau^-)$ would be negative.

By convention, case estimates should be strictly positive. This is enforced *a fortiori* by requiring that

$$\kappa y_{ir}(\tau^-) \geq c_{ir}(\tau^-), \tag{5}$$

for all i, r and $\tau < \tau_{ir}$, where $0 < \kappa < 1$ is a constant. Thus, if (2) or (3) yields a value of $y_{ir}(\tau^-)$ that breaches (5), then it is corrected to

$$y_{ir}(\tau^-) = \kappa^{-1} c_{ir}(\tau^-). \tag{6}$$

Constraint (5) has an equivalent form, that is useful in application to (3). This is

$$\begin{aligned} c_{ir}(\tau^-) \leq \kappa y_{ir}(\tau^-) &= \kappa (c_{ir}(\tau^-) + x_{ir}(\tau^-)) \\ &= \kappa \left(c_{ir}(\tau^-) + \frac{x_{ir}(\tau)}{g_{irl}^{Mi}} \right), \end{aligned}$$

from which follows

$$g_{irl}^{Mi} \leq \frac{\kappa}{1 - \kappa} \frac{x_{ir}(\tau)}{c_{ir}(\tau^-)}. \tag{7}$$

4.4.2 Base inflation adjustment

The allowance for inflation in case estimates is explained in Section 2.2. At each revision of incurred loss, an adjustment is made for the base inflation that has occurred since the previous revision. Within the simulator, these adjustments are made in reverse chronological order.

Consider an adjustment for the period between a revision at delay τ^* and $\tau^- > \tau^*$, where it is possible that $\tau^* = 0$, and also possible that $\tau = \tau_{ir}$. Let delay τ correspond to time \bar{t} , i.e. $\bar{t} = u_{ir} + v_{ir} + \tau$. Similarly, let delay τ^* correspond to time $\bar{t}^* = \bar{t} - (\tau - \tau^*)$. Then the adjustment factor (in reverse time) is $f(\bar{t}^*)/f(\bar{t})$.

As an example, if $y_{ir}(\tau^-)$ is computed by either (2) or (3), and the revision immediately prior to delay τ is a major revision at τ^* , then (2) at that epoch is replaced by

$$y_{ir}(\tau^*) = y_{ir}(\tau^-) \frac{f(\bar{t}^*)}{f(\bar{t})} = \frac{y_{ir}(\tau)}{g_{irl}^{Ma}} \frac{f(\bar{t}^*)}{f(\bar{t})}. \tag{8}$$

4.5 The treatment of “manual adjustments”

“Manual adjustments” occur twice within the default version of SPLICE, firstly in the requirement that minor and major revisions cannot occur simultaneously (Section 4.3.2), and secondly in the requirement of strict positivity of case estimates (Section 4.4.1).

It may happen that a user has cause to fit the simulation model to a real data set by maximum likelihood or by optimisation of some other statistical criterion. In this case, a pure mathematical statement of the model will be required, and any “manual adjustments” will require accommodation within it.

It should first be pointed out that the default model is quite optional, and the user is free to modify it in any form desired, eliminating “manual adjustments” if necessary. However, it can also be pointed out that the default model is in fact convertible to pure mathematical form without manual adjustments.

Consider, for example, the requirement that major and minor revisions not coincide. Coincidence could occur in the default implementation only if a minor and major revisions were simulated to coincide with the last major payment, where $\tau = \tau_{irl}^{Ma} = \omega_{ir}^{(m_{ir}-1)}$ (see Section 4.3.2). The df of the minor revision factor g_{irl}^{Mi} , according to Section 4.3.3, is $F_{G|\omega,\tau}^{Mi}$, in other words dependent on the epoch of the minor revision in question. The over-riding of the minor revision could be incorporated into this df by extension of the conditioning of $F_{G|\omega,\tau}^{Mi}$ to $F_{G|\omega,\tau,\tau_{irl}^{Ma},\omega_{ir}^{(m_{ir}-1)}}^{Mi}$ and stipulation within $F_{G|\omega,\tau,\tau_{irl}^{Ma},\omega_{ir}^{(m_{ir}-1)}}^{Mi}$ that $g_{irl}^{Mi} = 1$ if $\tau_{irl}^{Ma} = \omega_{ir}^{(m_{ir}-1)}$.

This is an extremely ugly expression of the model, and is not recommended for model description. It does illustrate, however, the way in which the over-riding of a minor revision by a major one can be expressed in the form of a genuine statistical model.

A similar device can be used to enforce the strict positivity of case estimates, as is achieved by (5)–(7). Again the device consists of expanding the conditionality of the random variable G_{irl}^{Mi} .

4.6 Aggregation of output

Consistent with SynthETIC, SPLICE provides both individual claim and aggregate output. In the latter, transactions are aggregated by accident and development period, where the duration of the periods may be chosen to any desired level of granularity (e.g. users who choose to work with calendar months can aggregate the transactions by month, quarter, or year). SPLICE by default uses accident and development quarters. The aggregate of case estimates for any particular development quarter includes the case estimates of all relevant individual claims at the **end of the quarter**. If a claim’s incurred loss is revised more than once during the quarter, only the estimated incurred loss after the last of those revisions will be reflected in the aggregate. Likewise, if the user chooses to summarise the claims on a yearly level, then the claim triangles will only capture the latest estimated incurred loss at the end of each year. As an illustration, Appendix B

shows the cumulative incurred loss triangle of the example implementation described in Section 5.1.

4.7 Out-of-bounds transactions

In this sub-section, a “transaction” includes occurrence, notification, settlement, a payment or a case estimate revision.

Sometimes, simulated transactions will take place beyond the end of the last development period. This out-of-bounds issue can potentially occur with all types of transaction (see also Avanzi *et al.*, 2021b). SPLICE treats those cases according to the following convention.

The bounds on development periods are ignored throughout that simulation, except in any aggregation of output by development period or addition of inflation, where any out-of-bounds transactions are counted as if at the end of the limiting development period I (where development periods are numbers 1, 2, . . .).

In the specific case of incurred losses, the simulated epoch of occurrence of any incurred loss revision is maintained throughout the simulation of details of the claim concerned, other than in the exceptions noted below. For example, if a minor revision occurs at development time $j > I$, and sizes of the minor revision multipliers depend on the epochs of the subject revisions, then the simulated value of j will be used in the simulation of those revision multipliers.

The epoch of revision is varied only at the stage where case estimates are assigned to development periods for the purpose of either tabulation or addition of inflation. In this context, the revision is assumed to have occurred at the end of development period I . In short, the integrity of epochs of transaction, and of any dependency on these epochs, is maintained throughout, with the sole exceptions of aggregation of out-of-bounds settlements and adjustment for inflation.

4.8 Data features

SPLICE inherits the claim payment structure from SynThETIC, which has been structured to resemble a real Auto Liability portfolio (“reference portfolio”). We refer to Section 4.3 of Avanzi *et al.* (2021b) for a review of the data features of the portfolio related to claim payments; see also Taylor *et al.* (2003), Taylor & McGuire (2004), McGuire (2007), McGuire *et al.* (2018).

As a result of the features alluded to in the previous paragraph, the portfolio behaviour could change over time with respect to claim payments (see Avanzi *et al.*, 2021b, for details). In contrast, the reference portfolio is not subject to time heterogeneity in the behaviour of its incurred loss estimates. There are, however, a few data features worthy of note. Some are briefly described in Principles 1–11. Broad details of one or two others are given immediately below. Full detail appears in Appendix A.

Principle 12. *The likelihood of a major revision at settlement increases with increasing claim size.*

Principle 13. *The timing of major revisions, other than those at settlement, is biased towards the early part of the claim’s lifetime.*

Principle 14. *The timing of minor revisions, other than those coincident with partial claim payments, is similarly biased.*

Although the generator of incurred loss estimates is time homogeneous, as noted just above, it does not follow that the behaviour of those estimates will be without complexity. As can be seen in Appendix A, the behaviour of the incurred loss estimates is dependent on that of the claim payments. The latter are time-heterogeneous, and some of the consequent complexity can be transmitted to the incurred loss estimates (see Sections 5.1.2 and 5.1.3).

5. Application of SPLICE

5.1 Example implementation of SPLICE with default parametrisation

5.1.1 Modular implementation in R

Avanzi *et al.* (2021b) performed an example simulation of claim payments in accordance with a detailed specification given there, with principal features of the experience similar to those of the reference portfolio. The generated experience covered 40 occurrence quarters, each tracked for 40 development quarters, with detailed transactional records.

SPLICE has been used to extend that simulation to include case estimates. Simulated paid losses remain unchanged from the earlier paper. The transactional simulation output now comprises key dates, and both claim payments and revisions of estimated incurred losses. A detailed specification of the modules involved in the simulation of case estimates is given in Appendix A.

Below we present the code to generate the example data set that is included as part of the package and described in the following Sections (5.1.2 and 5.1.3). We refer to the SPLICE vignette and package documentation for details of the function usage (Avanzi *et al.*, 2021a).

```

1  library(SPLICE)
2  set.seed(20201006)
3  test_claims <- SynthETIC::test_claims_object
4
5  # major revisions
6  major <- claim_majRev_freq(test_claims)
7  major <- claim_majRev_time(test_claims, major)
8  major <- claim_majRev_size(major)
9
10 # minor revisions
11 minor <- claim_minRev_freq(test_claims)
12 minor <- claim_minRev_time(test_claims, minor)
13 minor <- claim_minRev_size(test_claims, major, minor)
14
15 # development of case estimates
16 test <- claim_history(test_claims, major, minor)
17 test_inflated <- claim_history(
18   test_claims, major, minor,
19   base_inflation_vector = rep(1.02^1/4) - 1, times = 80))
20
21 # transactional data
22 test_incurred_dataset_noInf <- generate_incurred_dataset(test_claims,
23   test)
24 test_incurred_dataset_inflated <- generate_incurred_dataset(
25   test_claims, test_inflated)

```

An excerpt of the transactional data set, `test_incurred_dataset_noInf` generated from the above code, is included in Appendix C. Those results can easily be aggregated into triangles; see Appendix B.

Note that using the default set of parametrisation (which has been loosely calibrated to the reference portfolio) does not require the user to input further arguments. In cases where alternative sampling distributions or dependence structures are desired, they can be easily incorporated using the SPLICE framework described in Figure D.1.

The above is implementing the modules sequentially in a transparent way, but the dataset can also be directly generated with a single function `generate_data()`, for varying levels of complexity; see Section 5.2.

5.1.2 Comparison with chain ladder

The paid loss data simulated by SynthETIC was in substantial breach of the chain ladder assumptions (Avanzi *et al.*, 2021b). Section 6.2.1 of that paper demonstrated that those breaches translated into equally dramatic errors in the chain ladder's forecast of a loss reserve.

The destructive data features of the earlier paper related essentially to heterogeneity of the claim payment simulation model over time. For example, large (resp. small) claims were affected by small (resp. large) rates of payment period SI. This caused the observed profile of paid losses by development period to shift as one moved from one occurrence period to another.

In the present case, the incurred loss simulation model contains no such time heterogeneities, as can be seen in its definition in Appendix A. There is, nonetheless, the potential for some of the paid loss heterogeneities to induce incurred loss heterogeneities. There are at least a couple of ways in which this could occur, but the dominant one is described as follows.

First, note that claim size and delay to settlement are positively associated (Avanzi *et al.*, 2021b). Then recall the remark just above on SI, as it affects large and small claims, and note that the effect of this would be to steadily shorten the distribution of delay to settlement with advancing occurrence period. Next note that both major and minor revisions of incurred loss are triangularly distributed over a range of epochs determined by the delay to settlement. It follows that the distribution of these epochs would steadily shorten with advancing occurrence period.

This has been investigated as follows:

- (a) We aggregated the part of the simulated individual incurred data relating to quarters 1 to 40 into a 40×40 triangle.
- (b) We derived a forecast of ultimate incurred loss (and hence outstanding claims up to development quarter 40) by applying a chain ladder. Note that there is little claim activity beyond development quarter 40.
- (c) We compared this forecast with the “actual” amount of outstanding claims, simulated for payment quarters 41 to 79.

The predicted shortening of the distribution of epochs of incurred loss revisions is indeed empirically confirmed by Figure 3, which plots as a solid line, for each of the 40 occurrence periods, the proportion of the ultimate incurred loss recognised in paid losses and case estimates at the end of the 10th development period. The figure also shows as a dashed line the smoothed version of the plot with a moving 5-average as the smoother. It can be seen that the proportion of incurred loss recognised increases from about 70% to about 90% over the span of the occurrence periods.

The forecast results are set out in Table 3, where the chain ladder exhibits persistent over-estimation of outstanding liability (see also Appendix B). This is typical of the chain ladder in an environment of diminishing delay to recognition of incurred loss.

5.1.3 Intra-model dependencies

Appendix A provides details of the general structure of incurred loss development in respect of a specific claim. One point of note is that it does not follow a Markov process, i.e. a given claim transaction may depend on the entire history of the claim rather than on just its status at the point of the transaction in question. For example, if a claim experiences a large upward revision, other than at notification, then any subsequent revision is unlikely to be large.

The reason for the inclusion of this condition in the simulator may be illustrated by means of an example. Consider a claim that is initially estimated as of size \$50,000, and suppose that this estimate is subsequently revised to \$500,000, i.e. by a factor of 10. It is possible that a further substantial revision would occur, to \$750,000 say, but a further revision by a factor 10 is unlikely.

In fact, a negative association between the magnitudes of the first and second major revision factors (after notification) is evident from Appendix A. This is illustrated by Figure 4, which plots the major revision factors of 654 simulated claims from the data set described in Section 5.1.1 that

Table 3. Forecast claim costs based on synthetic data set.

Occurrence quarters	Estimated loss reserve		
	Target (Simulator)	Chain ladder estimate	Deviation of chain ladder from target
	\$M	\$M	%
1–10	9.2	11.4	23
11–20	26.2	32.3	23
21–25	26.1	35.7	37
26–30	61.8	74.9	21
31	20.5	23.3	13
32	24.4	27.9	15
33	27.8	25.2	– 9
34	30.1	35.7	19
35	26.4	33.9	28
36	36.2	43.2	19
37	33.0	40.5	23
38	43.1	65.9	53
39	39.0	43.5	12
40	46.2	29.1	–37
Total	449.6	522.6	16




Figure 3. Time heterogeneity of incurred losses over occurrence periods.

experience two major revisions (in addition to the one at notification), with a superimposed loess curve. A “revision factor”, as referred to here, is defined as the ratio by which the estimate of an individual claim’s incurred loss changes at any epoch. Inflationary effects have been excluded. The empirical correlation of the two major revision factors is estimated to be -0.617 for the example implementation.

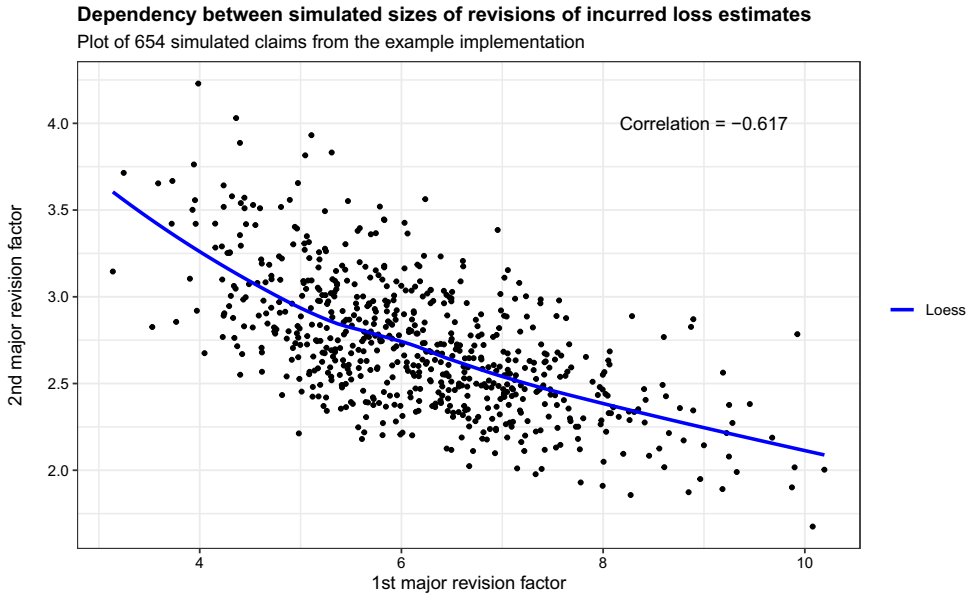


Figure 4. Dependency between simulated sizes of revisions of incurred loss estimates.

5.2 Alternative data sets of varying complexity

Section 5.1.2 demonstrates the failure of the chain-ladder model to capture the time heterogeneities in the incurred losses of the default portfolio. However, the flexibility of SPLICE ensures that the user can generate data of almost any level of complexity. Indeed, as discussed in Section 1.1, the user may generate a collection of data sets with varying levels of complexity and test the proposed model on the whole spectrum, in order to derive insights into its strengths and weaknesses.

5.2.1 Complexity scenarios

For the convenience of the user, we provide a data generation function in SPLICE that outputs alternative data sets under five hypothetical scenarios ranging in data complexity. The `generate_data()` function takes in a `complexity` parameter taking integer values between 1 (the simplest) and 5 (the most complex). Table 4 presents the detailed description of each scenario.

The most complex case (5) is the default illustrated and discussed in the previous section, whereas the simplest case (1) corresponds to a chain-ladder environment as described below. The intermediate cases allow focus on particular features, and to dial complexity up or down as required.

For the simplest case, chain-ladder compatibility is achieved in the following way:

- (a) all of SPLICE components Module 1 to Module 10 are defined to be independent of occurrence period;
- (b) base inflation and calendar period superimposed inflation in Modules 8 and 11 occur at a constant rate per period; and
- (c) occurrence period superimposed inflation in Modules 8 and 11 must be independent of all other components, but otherwise can be arbitrary.

These conditions were noted in Avanzi *et al.* (2021b) to be sufficient for homogeneity of paid losses over accident years. The time homogeneity of the incurred loss generation, noted in

Table 4. Descriptions of alternative data sets considered.

Complexity	1 (simplest)	2	3	4	5 (default, most complex)
Notification delay dependence structure (Module 3)	Independent of all features	Depends on claim size	Depends on claim size	Depends on claim size	Depends on claim size
Settlement delay dependence structure (Module 4)	Depends on claim size	Depends on claim size	Depends on claim size and occurrence period (steady decline over time)	Depends on claim size	Depends on claim size and occurrence period (decline over first 20 periods, plus a structural break at time 20 for smaller claims)
Base inflation (Module 8, 11)	Nil	Constant at 2% p.a.	Constant at 2% p.a.	Constant at 2% p.a.	Constant at 2% p.a.
Calendar period superimposed inflation (Module 8, 11)	Nil	Nil	Nil	Inflation shock at time 30 (from 0% to 10% p.a.)	Varies by claim size (up to 30% p.a. for the smallest claims, and 0 for claims exceeding \$200,000)
Occurrence period superimposed inflation (Module 8, 11)	Nil	Nil	Nil	Nil	Nil for the first 20 occurrence periods; for the next 20 periods, the size of the inflation varies by claim sizes (up to 40% reduction for the smallest claims)

Section 4.8 (Principle 14), will then ensure that the expected distribution of incurred loss estimates across development periods is the same for all occurrence periods before the inclusion of inflation. Just as for paid losses, occurrence period SI is directly reflected in the chain ladder row parameters, which can be arbitrary. To consider the other forms of inflation, recall from the discussion in Section 2.2 that case estimates of incurred loss include base inflation to the date of estimate and calendar period SI up to settlement date.

It is already noted above that the expected distribution of incurred loss estimates across development periods is the same for all occurrence periods before the inclusion of inflation, and this will remain the case when base inflation and calendar period SI are added at a constant (though possibly different) rates.

The above demonstrates that SPLICE allows the user to construct very simple or very complex data sets. Naturally, a large array of intermediate cases sit in between. Full control of that level of complexity (and knowledge of its origin) allows modelers to test or illustrate specific strengths or weaknesses of their contending model, as discussed in Section 1. This idea is used, for instance, in McGuire *et al.* (2018) and Al-Mudafer *et al.* (2021), with SynthETIC.

We remark that the `generate_data()` function by no means defines the limits of the complexity that can be achieved with SPLICE. The function is provided for the convenience of users who wish to generate (a collection of) data sets under some representative scenarios. If more complex features are required, the user is free to modify the distributional assumptions to achieve their purpose.

5.2.2 Comparison with chain ladder

Figure 5 plots the distribution of the chain ladder estimation error for the five scenarios described above, with 500 simulations for each scenario. The estimation error here refers to the percentage deviation of the chain ladder estimate of outstanding claim liabilities from the true value (i.e. produced by the simulator).

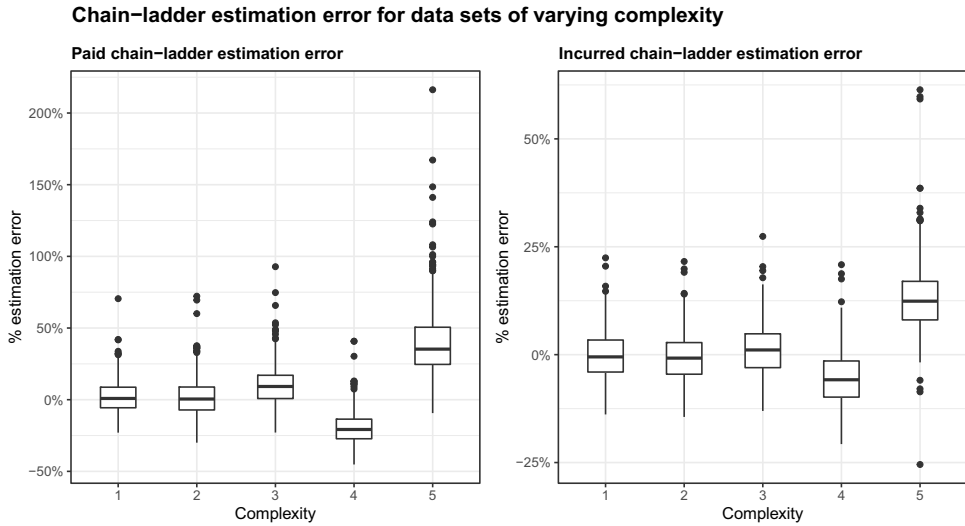


Figure 5. Chain-ladder estimation error for the five different complexity scenarios, with 500 simulations each.

The plot demonstrates that, regardless of whether we use paid (left panel) or incurred triangles (right panel) to estimate reserves, there is a consistent increase in the chain ladder estimation error as the complexity of the data set grows. This is an immediate result of the design of the data sets as outlined in Section 5.2.1: the simplest scenario is designed to provide chain ladder compatibility, whereas the most complex scenario is in complete breach of chain ladder assumptions, as visualised in Appendix D.

A side observation is that the incurred chain ladder method (right panel) consistently produces significantly lower prediction errors than the paid method, which provides a further motivation for a simulator like SPLICE that is capable of generating case estimates. The relatively large error associated with the paid method is driven by the unstable age-to-age factors in the earlier development periods, which translate into highly variable estimates of ultimate claims for the most recent occurrence periods.

The plot illustrates the range of complexity that can be achieved with SPLICE. A modeler interested in testing the effectiveness of their model in dealing with different scenarios can use the above as a starting point.

6 Conclusion

SPLICE is a CRAN claim simulator package that extends *Synthetic* (Avanzi *et al.*, 2021c) so that it now simulates both claim payments and estimates of incurred loss, and this enhanced version is described in the foregoing sections. Its users can access and modify its code freely, and it comes with a full set of default options which have been designed to be realistic as described in Principles 1–14, as well as [4.3.1] to [4.3.7] of Avanzi *et al.* (2021b).

Of previously existing simulators, to the extent of the authors' knowledge, only one simulator is capable of generating sequence of case estimates of incurred losses through the lifetime of each claim. As outlined in Section 1.2.1, it is subject to several limitations, many of which are addressed in SPLICE. The development of such a data generator is motivated by the scarcity of granular data with realistic features, the availability of which is essential for model development (Embrechts & Wüthrich, 2022).

The backward computation algorithm that SPLICE adopts guarantees that the final incurred estimate coincides with the claim size. This simplifies the task of generating the evolution history

of case estimates to the simulation of major and minor revisions. SPLICE can flexibly generate the different variates of major and minor revisions. Furthermore, thanks to its modular structure, SPLICE allows the user to modify its default dependence structure by adjusting parameters within a module or replacing with their own. This enables the user to validate a proposed model against data of any desired level of complexity, as facilitated by the function `generate_data()` and discussed in Section 5.2.

By default, SPLICE incorporates complex dependencies that reflect a realistic claim process (e.g. between different revisions of incurred loss estimates), producing desirable data features outlined in Section 4.8. This may be of use in testing granular models, which sometimes include unrealistic assumptions of independence between different components. SPLICE may be of especial value in testing models that estimate reserves from incurred (and paid) loss data, such as the paid incurred chain reserving method of Merz & Wüthrich (2010).

Acknowledgements. The authors gratefully acknowledge productive discussions on the extension of SynthETIC with Bernard Wong. The authors are also grateful to William Ho for research assistance, and comments that led to improvements of the R code. This research was supported under Australian Research Council's Linkage (LP130100723) and Discovery (DP200101859) Projects funding schemes. Melantha Wang acknowledges financial support from UNSW Australia Business School. The views expressed herein are those of the authors and are not necessarily those of the supporting organisations.

References

- Al-Mudafer, M.T., Avanzi, B., Taylor, G.C. & Wong, B. (2021). Stochastic loss reserving with mixture density neural networks. *Insurance: Mathematics and Economics*, **105**, 144–174.
- Avanzi, B., Taylor, G. & Wang, M. (2021a). SPLICE: synthetic paid loss and incurred cost experience (splice) simulator. Available online at the address <https://CRAN.R-project.org/package=SPLICE>.
- Avanzi, B., Taylor, G., Wang, M. & Wong, B. (2021b). SynthETIC: an individual insurance claim simulator with feature control. *Insurance: Mathematics and Economics*, **100**, 296–308.
- Avanzi, B., Taylor, G., Wang, M. & Wong, B. (2021c). SynthETIC: synthetic experience tracking insurance claims. Available online at the address <https://CRAN.R-project.org/package=SynthETIC>.
- Bear, R., Shang, K. & You, H. (2020). cascsim: casualty actuarial society individual claim simulator. Available online at the address <https://CRAN.R-project.org/package=cascsim>.
- CAS Loss Simulation Model Working Party (2007). Parameterizing the loss simulation model. Available online at the address <https://www.casact.org/research/lsmwp/bsupaper.pdf>.
- CAS Loss Simulation Model Working Party (2011). Modeling loss emergence and settlement processes. Casualty Actuarial Society E-Forum.
- Coté, M.-P., Hartman, B., Mercier, O., Meyers, J., Cummings, J. & Harmon, E. (2020). Synthesizing property & casualty ratemaking datasets using generative adversarial networks.
- De Felice, M. & Moriconi, F. (2019). Claim watching and individual claims reserving using classification and regression trees. *Risks*, **7**(4), 1–36.
- Dutang, C. & Charpentier, A. (2019). CASdatasets: insurance datasets. Available online at the address <http://cas.uqam.ca/>.
- Dutang, C., Goulet, V. & Pigeon, M. (2008). actuar: an R package for actuarial science. *Journal of Statistical Software*, **25**(7), 38.
- Embrechts, P. & Wüthrich, M.V. (2022). Recent challenges in actuarial science. *Annual Review of Statistics and its Application*, **9**, 119–140.
- Gabrielli, A. & Wüthrich, M.V. (2018). Individual claims history simulation machine. *Risks*, **6**(2), 29.
- Harej, B., Gächter, R. & Jamal, S. (2017). Individual claim development with machine learning. Available online at the address http://www.actuaries.org/ASTIN/Documents/ASTIN_ICDML_WP_Report_final.pdf.
- Mack, T. & Quarg, G. (2004). Munich Chain Ladder. *Blätter der DGVFM*, **26**(4), 597–630.
- McGuire, G. (2007). Individual claim modelling of CTP data. In *XIth Accident Compensation Seminar*. Institute of Actuaries of Australia, Sydney, Australia.
- McGuire, G., Taylor, G. & Miller, H. (2018). Self-assembling insurance claim models using regularized regression and machine learning. Variance (in press). Also in SSRN.
- Merz, M. & Wüthrich, M.V. (2010). Paid-incurred chain claims reserving method. *Insurance: Mathematics and Economics*, **46**(3), 568–579.
- Taylor, G. (2019). Claim models: granular and machine learning forms. *Risks*, **7**(3), 82.

Taylor, G. & McGuire, G. (2004). Loss reserving with glms: a case study. Casualty Actuarial Society Discussion Paper Program, Applying and Evaluating Generalised Linear Models.
 Taylor, G., McGuire, G. & Greenfield, A. (2003). Loss reserving: past, present and future, University of Melbourne Research Paper.
 Taylor, G., McGuire, G. & Sullivan, J. (2008). Individual claim loss reserving conditioned by case estimates. *Annals of Actuarial Science*, 3(1–2), 215–256.
 Teugels, J.L. & Sundt, B. (Eds.) (2004). *Encyclopedia of Actuarial Science*. Chichester, UK: John Wiley & Sons, Ltd.

Appendix A. Parametrisations

The following table displays the formal parameterisation of Modules 9–11 for the example of Section 5.1.

Parameter type	Functional form	Numerical parameters
Global		Time unit = 1/4 (Calendar quarter) Reference claim size = 200, 000
Major revisions		
Number of revisions of incurred loss	$F_{K s}^{Ma}(k; s)$	For $s_{ir} \leq 15,000$, $\mathbb{P}(K_{ir}^{Ma} > 1) = 0$ For $s_{ir} > 15,000$: $\mathbb{P}(K_{ir}^{Ma} = 2) = 0.1 + 0.3 \min(1, \frac{s_{ir}-15,000}{185,000})$; $\mathbb{P}(K_{ir}^{Ma} = 3) = 0.5 \min(1, \frac{\max(0, s_{ir}-15,000)}{185,000})$; $\mathbb{P}(K_{ir}^{Ma} > 3) = 0$
Epochs of revisions of incurred loss	$\mathbb{P}(\tau_{irl}^{Ma} = w_{ir}^{(m_{ir}-1)} s_{ir})$ for $l = k_{ir}^{Ma}$	Major revisions occur only for $m_{ir} \geq 4$ $\mathbb{P}(\tau_{irl}^{Ma} = w_{ir}^{(m_{ir}-1)} s_{ir}) = 0.2 \min(1, \frac{\max(0, s_{ir}-200,000)}{2,800,000})$
	$F_{\tau w}^{Ma1}(\tau; w)$	The τ_{irl}^{Ma} , $l = 2, \dots, k_{ir}^{Ma} - 1$ are sampled from a triangular pdf with range $(w_{ir}^{(m_{ir}-1)}/3, w_{ir}^{(m_{ir}-1)})$ and maximum density at $w_{ir}^{(m_{ir}-1)}/3$
	$F_{\tau w}^{Ma2}(\tau; w)$	The τ_{irl}^{Ma} , $l = 2, \dots, k_{ir}^{Ma}$ are sampled from a triangular pdf with range $(w_{ir}/3, w_{ir})$ and maximum density at $w_{ir}/3$
Revision factors (multipliers of incurred losses)	$F_{G h}^{Ma}(g; h)$	For $l = 2$, $\ln g_{irl}^{Ma}$ drawn from $\mathcal{N}(1.8, 0.2^2)$ For $l = 3$, $\ln g_{irl}^{Ma}$ drawn from $\mathcal{N}(\mu, 0.2^2)$ with $\mu = 1 + 0.07(6 - g_{ir2}^{Ma})$ In each case, g_{irl}^{Ma} is subject to an upper bound of $\frac{0.95y_{ir}(\tau)}{c_{ir}(\tau)}$, where the revision is assumed to occur with delay τ from notification
Minor revisions		
Number of revisions of incurred loss		
Simultaneous with partial payment	$F_{B l}(b)$	$F_{B l}(0) = F_{B l}(1) = \frac{1}{2}$
Not simultaneous with partial payment	$F_{K w}^{Mi2}(k; w)$	Distribution of K_{ir}^{Mi2} is geometric, with mean = $\min(3, \frac{w_{ir}}{4})$
Epochs of revisions of incurred loss		

Parameter type	Functional form	Numerical parameters
Simultaneous with partial payment		Coincide with the partial payments
Not simultaneous with partial payment	$F_{T w}^{Mi}(\tau; w)$	The $\tau_{irl}^{Mi}, l = 1, \dots, k_{ir}^{Mi}$ are drawn from a uniform distribution with range $(w_{ir}/6, w_{ir})$
Revision factors (multiplier of outstanding paid losses)	$F_{G w,\tau}^{Mi}(g; w, \tau)$	For $\tau_{irl}^{Mi} \leq \frac{w_{ir}}{3}$, $\ln g_{irl}^{Mi}$ is drawn from: $\mathcal{N}(0.15, 0.05^2)$ if preceded by a 2 nd major revision (i.e. subsequent to notification); $\mathcal{N}(0.15, 0.1^2)$ otherwise For $\frac{w_{ir}}{3} < \tau_{irl}^{Mi} \leq \frac{2w_{ir}}{3}$, $\ln g_{irl}^{Mi}$ is drawn from: $\mathcal{N}(0, 0.05^2)$ if preceded by a 2 nd major revision; $\mathcal{N}(0, 0.1^2)$ otherwise For $\tau_{irl}^{Mi} < \frac{2w_{ir}}{3}$, $\ln g_{irl}^{Mi}$ is drawn from: $\mathcal{N}(-0.1, 0.05^2)$ if preceded by a 2 nd major revision; $\mathcal{N}(-0.1, 0.1^2)$ otherwise In each case, g_{irl}^{Mi} is subject to the constraint that $y_{ir}(\tau^-) = c_{ir}(\tau^-) + x_{ir}(\tau^-) \geq k c_{ir}(\tau^-)$, or equivalently $x_{ir}(\tau^-) = (y_{ir}(\tau^-) - c_{ir}(\tau^-))/g_{irl}^{Mi} \geq (k - 1)c_{ir}(\tau^-)$, where the revision is assumed to occur with delay τ from notification and $k > 1$ is a fixed constant (not necessarily the same k as above). The current program assumes $1/k = 0.95$
Base inflation	$f(\bar{t})$	$= (1 + \alpha)^{\bar{t}}$, where α is equivalent to an increase of 2% p.a., allowing for the relevant time units
Superimposed inflation	$g_O(u s)$	$= 1$ for $u \leq 20$ $= 1 - 0.4 \max\left(0, 1 - \frac{s}{50,000}\right)$ for $u > 20$
	$g_C(\bar{t} s)$	$= (1 + \beta(s))^{\bar{t}}$ with $\beta(s) = \gamma \max\left(0, 1 - \frac{s}{200,000}\right)$ and γ is equivalent to a 30% p.a. inflation rate, allowing for the relevant time units

Notes about claim sizes and inflation:

(a) Some components are defined in terms of claim size. The definition then displays claim size in raw uninflated units. The inputs to the example application of SPLICE, on the other hand, express claim sizes as multiples of a reference claim size equal to 200,000. For example, the amount of 15,000 that appears in the definition of the frequency of major revisions is expressed in SPLICE as $0.075 \times 200,000$.

(b) Wherever claim size (s or s_{ir}) is compared with a numerical quantity (e.g. 50,000 for $S1$), the claim size excludes all forms of inflation.

Appendix B. Cumulative Incurred Loss Triangle of the Example Implementation

For space considerations, below we show only the claim development triangles on a yearly basis (with the inclusion of inflation); however, the underlying data is calculated based on quarterly development pattern and is available on the SPLICE repository (see Section 1.3).

The code to produce the triangles in Tables B.1 and B.2 builds on the sample code in Section 5.1.1 and is provided below:

```

1  ## SPLICE simulated incurred loss triangle
2  incurred_inflated <- output_incurred(
3    test_inflated, incremental = F, aggregate_level = 4)
4
5  ## Chain-ladder predicted incurred loss
6  # output the past cumulative incurred triangle simulated by SPLICE
7  cumtri <- output_incurred(
8    test_inflated, aggregate_level = 4, incremental = F, future = F)
9  # calculate the age to age factors
10 selected <- attr(ChainLadder::ata(cumtri), "vwtd")
11 # complete the triangle
12 CL_prediction <- cumtri
13 J <- nrow(cumtri)

```

```

14 for (i in 2:J) {
15   for (j in (J - i + 2):J) {
16     CL_prediction[i, j] <- CL_prediction[i, j - 1] * selected[j - 1]
17   }
18 }

```

Table B.1. Cumulative incurred loss triangle of simulated claims (\$000).

Accident Year	Development Year									
	1	2	3	4	5	6	7	8	9	10
1	16,389	39,274	52,490	64,217	66,874	73,648	76,208	77,264	77,173	77,211
2	18,017	42,795	58,207	67,869	74,310	75,871	78,755	80,789	80,997	81,360
3	21,714	52,799	65,190	71,618	78,491	81,707	82,539	83,320	83,431	83,230
4	21,172	48,529	64,916	75,795	83,538	87,737	87,446	87,523	87,610	86,801
5	26,459	64,570	80,520	94,317	98,718	100,446	102,592	103,159	102,880	103,254
6	29,379	71,254	90,285	100,748	104,567	106,252	108,053	107,613	107,433	107,352
7	21,154	64,467	82,384	97,824	103,947	107,824	108,037	107,518	107,332	107,334
8	38,244	82,992	104,322	114,536	117,436	120,199	122,193	124,254	125,922	125,889
9	34,735	97,448	125,597	140,065	146,810	148,883	151,319	152,575	152,411	152,916
10	49,339	120,973	149,794	160,851	167,556	167,292	166,790	166,531	166,428	166,216

Table B.2. Cumulative claim development triangle predicted by the chain ladder model (\$000).

Accident Year	Development Year									
	1	2	3	4	5	6	7	8	9	10
1	16,389	39,274	52,490	64,217	66,874	73,648	76,208	77,264	77,173	77,211
2	18,017	42,795	58,207	67,869	74,310	75,871	78,755	80,789	80,997	81,037
3	21,714	52,799	65,190	71,618	78,491	81,707	82,539	83,320	83,382	83,424
4	21,172	48,529	64,916	75,795	83,538	87,737	87,446	88,871	88,938	88,982
5	26,459	64,570	80,520	94,317	98,718	100,446	102,331	103,998	104,076	104,127
6	29,379	71,254	90,285	100,748	104,567	109,114	111,162	112,973	113,058	113,113
7	21,154	64,467	82,384	97,824	104,406	108,946	110,991	112,800	112,884	112,940
8	38,244	82,992	104,322	120,878	129,012	134,622	137,148	139,383	139,487	139,556
9	34,735	97,448	124,934	144,761	154,502	161,221	164,246	166,923	167,048	167,130
10	49,339	122,472	157,017	181,935	194,177	202,621	206,423	209,787	209,944	210,048

Appendix C. Excerpt of Simulated Data

Table C.1 is an excerpt of the transactional dataset generated by SPLICE (prior to the addition of inflation) and displays the full claim history of claims #2 and #40 (in Figure 2) from notification to settlement. Table C.2 provides the detailed description of the variables.

Equations (1)–(6), which describe the computation of case estimates before the effect of inflation is added, may be illustrated by a numerical example. The third transaction record for claim #2 indicates a simultaneous minor revision at the time of a partial payment, by a factor of 1.0503. The total incurred loss after this transaction is thus the sum of the revised outstanding paid loss

Table C.1. Transaction records of two sample claims from the example implementation.

claim_no	claim_size	txn_time	txn_delay	txn_type	incurred	OCL	cumpaid	multiplier
2	22,562	1.298	0.000	Ma	21,635	21,635	0	1.0000
2	22,562	2.203	0.906	PMi	23,694	21,688	2,005	1.0952
2	22,562	2.695	1.397	PMi	24,784	20,654	4,130	1.0503
2	22,562	3.317	2.019	PMi	22,562	2,446	20,116	0.8924
2	22,562	3.629	2.332	P	22,562	0	22,562	
40	143,183	4.278	0.000	Ma	15,969	15,969	0	1.0000
40	143,183	4.456	0.177	P	15,969	13,901	2,068	
40	143,183	6.062	1.784	P	15,969	11,957	4,013	
40	143,183	7.335	3.056	P	15,969	9,296	6,674	
40	143,183	7.503	3.225	Mi	17,402	10,728	6,674	1.1541
40	143,183	7.537	3.258	Mi	18,318	11,644	6,674	1.0854
40	143,183	8.177	3.899	P	18,318	9,846	8,472	
40	143,183	9.116	4.837	Mi	19,218	10,745	8,472	1.0913
40	143,183	9.121	4.842	P	19,218	8,972	10,245	
40	143,183	10.162	5.884	P	19,218	6,929	12,289	
40	143,183	10.480	6.202	P	19,218	5,141	14,077	
40	143,183	11.133	6.855	PMi	18,769	2,891	15,878	0.9127
40	143,183	12.017	7.738	Mi	18,918	3,040	15,878	1.0516
40	143,183	12.442	8.163	PMi	18,595	930	17,665	0.8938
40	143,183	12.495	8.216	Mi	18,595	930	17,665	1.0149
40	143,183	12.568	8.290	Ma	53,754	36,089	17,665	6.1785
40	143,183	13.623	9.345	Mi	54,116	36,451	17,665	1.0100
40	143,183	13.697	9.418	PMi	52,969	33,242	19,727	0.9685
40	143,183	14.102	9.824	Ma	168,224	148,497	19,727	3.1759
40	143,183	14.109	9.830	P	168,224	146,744	21,480	
40	143,183	14.925	10.646	Mi	164,702	143,222	21,480	0.9760
40	143,183	15.402	11.123	Mi	149,735	128,255	21,480	0.8955
40	143,183	15.589	11.310	P	149,735	126,591	23,145	
40	143,183	17.114	12.836	PMi	147,069	17,189	129,880	0.9789
40	143,183	17.414	13.136	Mi	145,194	15,314	129,880	0.8909
40	143,183	18.334	14.055	PMi	143,183	0	143,183	0.8687

(21,688 \times 1.0503) and the cumulative paid loss prior to the revision (\$2,005), as governed by equation (4). The partial payment made brings the outstanding claim liability down to \$20,654.

Major revisions are rarer, but usually of a greater magnitude. Claim #40 sees a major revision 9.824 periods after notification (at transaction time 14.102). The major revision effects a change on the incurred loss directly (52,969 \times 3.1759), driving up the estimated incurred loss and outstanding payments to \$168,224 and \$148,497, respectively.

We remark that in practice, the computation algorithm simulates the case estimates in reverse chronological order to ensure that the final incurred estimate coincides with the claim size; see Section 4.

The code to generate this data set is provided in Section 5.1.1.

Table C.2. Variable descriptions.

Variable	Description
claim_no	claim number, which uniquely characterises each claim
claim_size	size of the claim (in constant dollar values)
txn_time	the calendar time of transaction (on a continuous time scale, with origin at the beginning of occurrence period 1)
txn_delay	delay from notification to the subject transaction
txn_type	nature of the transactions, “Ma” for major revision, “Mi” for minor revision, “P” for payment, “PMa” for major revision coincident with a payment, “PMi” for minor revision coincident with a payment
incurred	case estimate of total incurred loss immediately after the transaction
OCL	case estimate of outstanding claim payments immediately after the transaction
cumpaid	cumulative claim paid after the transaction
multiplier	revision multipliers (subject to further constraints specified by equations (5) and (6)), missing for transactions that do not involve a revision

Appendix D. Visualisations of Claim Development under Different Complexity Scenarios

For space considerations, we show only the plots for the complexity level of 1 and 5. At a complexity level of 1, the clustered lines across all occurrence periods in both plots indicate the homogeneity of claim payments and incurred losses across all occurrence periods. This contrasts with the two plots in the bottom featuring a steady shortening of both paid and incurred patterns with respect to the occurrence period, which makes the modeling a challenge.

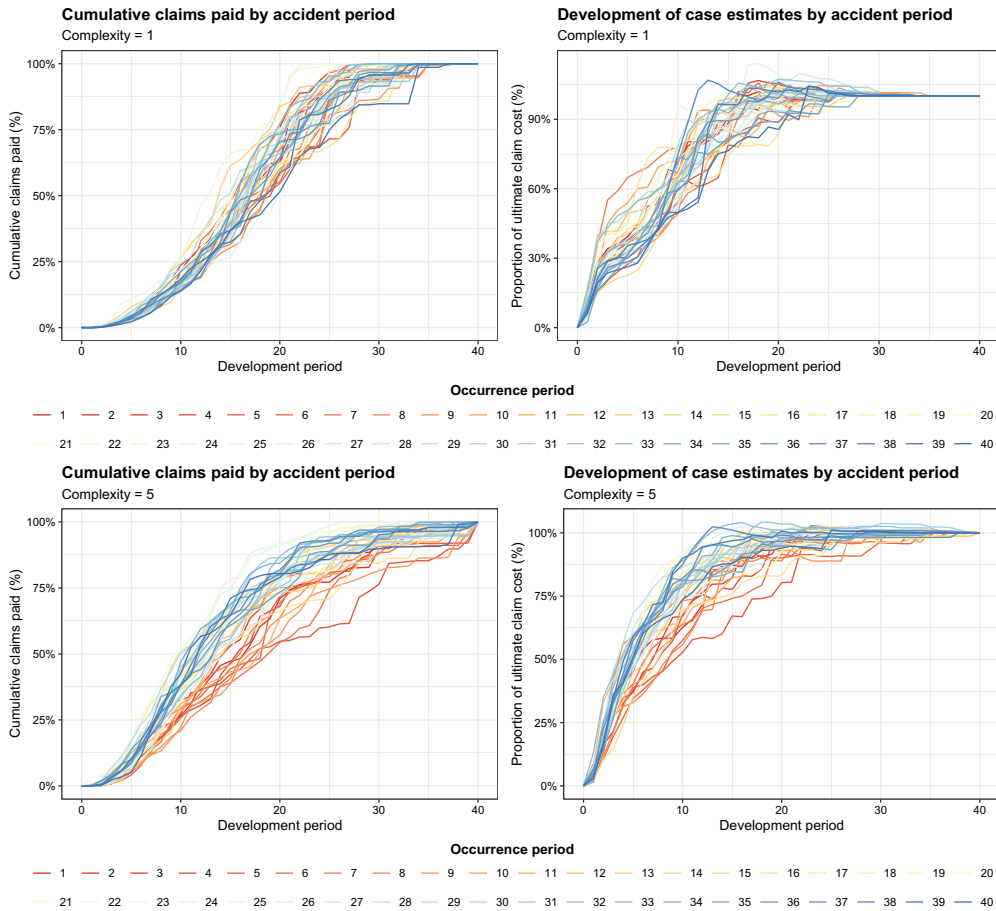


Figure D.1. Claim development pattern under different complexity scenarios (1 represents the simplest, 5 represents the most complex, in our case, the default portfolio described in Section 5.1).