



Minerva Access is the Institutional Repository of The University of Melbourne

Author/s:

Avanzi, B;Taylor, G;Wang, M;Wong, B

Title:

SynthETIC: An individual insurance claim simulator with feature control

Date:

2021-07-07

Citation:

Avanzi, B., Taylor, G., Wang, M. & Wong, B. (2021). SynthETIC: An individual insurance claim simulator with feature control. *INSURANCE MATHEMATICS & ECONOMICS*, 100, pp.296-308. <https://doi.org/10.1016/j.insmatheco.2021.06.004>.

Persistent Link:

<https://hdl.handle.net/11343/324127>

# SynthETIC: an individual insurance claim simulator with feature control

Benjamin Avanzi<sup>a</sup>, Greg Taylor<sup>\*,b</sup>, Melantha Wang<sup>b</sup>, Bernard Wong<sup>b</sup>

<sup>a</sup>Centre for Actuarial Studies, Department of Economics, University of Melbourne VIC 3010, Australia  
<sup>b</sup>School of Risk and Actuarial Studies, UNSW Australia Business School, UNSW Sydney NSW 2052, Australia

---

## Abstract

Recent years have seen rapid increase in the application of machine learning to insurance loss reserving. They yield most value when applied to large data sets, such as individual claims, or large claim triangles. In short, they are likely to be useful in the analysis of any data set whose volume is sufficient to obscure a naked-eye view of its features. Unfortunately, such large data sets are in short supply in the actuarial literature. Accordingly, one needs to turn to synthetic data. Although the ultimate objective of these methods is application to real data, the use of synthetic data containing features commonly observed in real data is also to be encouraged.

While there are a number of claims simulators in existence, each valuable within its own context, the inclusion of a number of desirable (but complicated) data features requires further development. Accordingly, in this paper we review those desirable features, and propose a new simulator of individual claim experience called **SynthETIC**.

Our simulator is publicly available, open source, and fills a gap in the non-life actuarial toolkit. The simulator specifically allows for desirable (but optionally complicated) data features typically occurring in practice, such as variations in rates of settlements and development patterns; as with superimposed inflation, and various discontinuities, and also enables various dependencies between variables. The user has full control of the mechanics of the evolution of an individual claim. As a result, the complexity of the data set generated (meaning the level of difficulty of analysis) may be dialled anywhere from extremely simple to extremely complex. The default version is parameterized so as to include a broad (though not numerically precise) resemblance to the major features of experience of a specific (but anonymous) Auto Bodily Injury portfolio, but the general structure is suitable for most lines of business, with some amendment of modules.

*Key words:* granular models, individual claims, individual claim simulator, loss reserving, partial payments, simulated losses, superimposed inflation, SynthETIC

JEL codes: C51, C52, C53, C55

MSC classes: 91G70, 91G60, 62P05

---

## 1. Introduction

Recent years have seen rapid increase in the application of **machine learning (ML)** to insurance loss reserving. Examples of various ML techniques applied in aggregated loss reserving include, for example, England and Verrall (2001) who applied GAM, Clemente et al. (2014) who considered GAMLSS, McGuire et al. (2018) who applied LASSO techniques. Neural-network based approaches were also considered by Wüthrich (2018b), Kuo (2019), and Gabrielli (2020). Similarly, ML techniques have also been used in models of individual claims reserving, including for example Wüthrich (2018a) who applied tree based

---

\*Corresponding author.

*Email addresses:* [b.avanzi@unimelb.edu.au](mailto:b.avanzi@unimelb.edu.au) (Benjamin Avanzi), [gregory.taylor@unsw.edu.au](mailto:gregory.taylor@unsw.edu.au) (Greg Taylor), [chenyi.wang@student.unsw.edu.au](mailto:chenyi.wang@student.unsw.edu.au) (Melantha Wang), [bernard.wong@unsw.edu.au](mailto:bernard.wong@unsw.edu.au) (Bernard Wong)

methods, and Mulquiney (2006), DeLong et al. (2020) who applied neural-network based models. See also Richman (2020a,b), and Blier-Wong et al. (2021) for additional commentary and extensive references.

These ML methods are hungry for data. They yield most value when applied to large data sets, such as individual claim, even transactional, data sets. They may also be useful in the analysis of large claim triangles, e.g.  $40 \times 40$ . In short, they are likely to be useful in the analysis of any data set whose volume is sufficient to obscure a naked-eye view of its features.

Unfortunately, such large data sets are in short supply in the actuarial literature, and so there is a shortfall in the material required to test ML methods. Accordingly, one needs to turn to synthetic data. Although the ultimate objective of these methods is application to real data, the use of synthetic data containing features commonly observed in real data is also to be encouraged.

Indeed, there is a school of thought that proposed claim models should **always** be tested against synthetic data, as well as real data if available. The reason is that the features contained in the synthetic data set will be known, by construction. One is then able to test the extent to which the proposed model is successful in identifying them.

There are several individual claim simulators in existence, and they are briefly reviewed in Section 5.1. Although each of these is valuable within its own context, we believe that there is scope for a further simulator due to the need for inclusion of a number of desirable (but complicated) data features, see Section 5. It is against this background that the individual claim simulator **SynthETIC** (**S**ynthetic **E**xperience **T**racking **I**nsurance **C**laims) is introduced. It is publicly available, open source (see Section 8), and is the only claims simulator available as a full R package on CRAN, with all the associated documentation.. Furthermore, it is fully compatible with existing actuarial CRAN packages such as **actuar** (Dutang et al., 2008), and fills a gap in the non-life actuarial toolkit.

**SynthETIC** simulates, for each claim, occurrence, notification, the timing and magnitude of individual partial payments, and closure. While **SynthETIC** comes with a default set of distributional assumptions, the flexibility of **SynthETIC** implies that users of the program are free to choose and implement any form of distribution models with ease. Inflation, including (optionally) superimposed inflation, is incorporated in payments. Superimposed inflation may occur over calendar or accident periods. The micro-level claim data can be summarized by accident and payment “periods” whose duration is an arbitrary choice (e.g. month, quarter, etc.) available to the user, which is then ready for chain ladder reserving analysis by **ChainLadder** (Gesmann et al., 2020).

The code is structured as eight modules (occurrence, notification, ...), any one or more of which may be re-designed according to the user’s requirements. The default version is parameterized so as to include a broad (though not numerically precise) resemblance to the major features of experience of a specific (but anonymous) Auto Bodily Injury portfolio. It thus reflects a number of desirable (but complicated) data features, but the general structure is suitable for most lines of business, with some amendment of modules. The structure of the simulator enables the inclusion of a number of important dependencies between the variables related to an individual claim, e.g. dependence of notification delay on claim size, of the size of a partial payment on the sizes of those preceding, etc. Many users should be able to use **SynthETIC** without modifying such dependence structure between individual claim features. However, **SynthETIC** is flexible enough to accommodate alternative dependencies. Ultimately, the user has full control of the mechanics of the evolution of an individual claim.

Thanks to the high flexibility of **SynthETIC**, the complexity of the data set generated (meaning the level of difficulty of analysis) *may be dialled anywhere from extremely simple to extremely complex*. At the extremely simple end would be chain-ladder-compatible data, and so the alternative data structures available enable proposed loss reserving models to be tested against more challenging data sets. Indeed, the user may generate a collection of data sets that provide a spectrum of complexity. This is particularly useful when trying to ‘break’ a model and ascertaining its qualities and limitations.

This paper is structured as follows. After a discussion of desirable features in Section 5, we present the notation and details of **SynthETIC** in the following sections. The claim process is defined by 8 modules (see Section 4.1), any of which may be unplugged and replaced by an alternative module of the user’s choosing. The current default parameterization broadly resembles a specific real Auto Liability data set that contains a number of features that render its modelling demanding.

After some notation in Section 3, the architecture of **SynthETIC** is described in Section 4. Section 5 discusses its relation to prior literature, while Section 6 gives detail of the default parameterization just mentioned. Section 7 discusses the circumstances and manner in which **SynthETIC** might be applied, Section 8 gives the web address of the code and the example data set used in Section 6, and Section 9 contains some closing comments.

## 2. Desirable data features

Much of the actuarial literature on loss reserving is focused on the chain ladder. This model requires a very specific parametric structure in which the expectation of the target variable is equal to the product of a row parameter and a column parameter (see, for example, Taylor, 2000; Wüthrich and Merz, 2008).

However, the same literature is dotted with examples of data sets that do not fit this prescription. Indeed, various models within the literature have been devised with the express purpose of addressing such non-conforming data sets.

This failure to conform with the chain ladder structure can arise in a number of ways. First, the **rate of settlement of claims** can change from one accident period to another. Models that take account of this feature have a long pedigree, e.g. Fisher and Lange (1973); Reid (1978) through to the present day, e.g. McGuire et al. (2018).

More generally, **development patterns** may change from one accident period to another. This may be due to changing rates of settlement, changing mixture of risks, or other causes. Again, models that seek to address this feature have a long pedigree, e.g. Berquist and Sherman (1977) up to recent times, e.g. Meyers (2015).

**Superimposed inflation** is a further potential disruptor of claim experience. The chain ladder model can accommodate inflation that occurs at a constant rate over time, but not variable-rate inflation. Examples of real data sets that appear to exhibit variable rates of superimposed inflation appear in e.g. Taylor (2000) and McGuire et al. (2018).

A claim triangle may experience a **discontinuity** as one passes from one accident period to another. This might result from a change in legislation governing claim conditions, or a sudden change in the type of business underwritten. An example is discussed by McGuire (2007).

Models that have been created to address data features such as the above require data sets containing those features if they are to be tested effectively. A major objective of **SynthETIC** is to provide this facility, and to our knowledge, it is the only simulator (at the time of writing this paper) that includes all those features natively.

## 3. Notation

**SynthETIC** works with **exact transaction times**, so time will be measured continuously. Calendar time  $\tau = 0$  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 4.1), 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  $1, 2, \dots$ , 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 (i.e. the final payment period corresponds to the final development period of the final occurrence period).

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 calendar time  $\tau = 0$ .

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 superimposed inflation occurs in two sub-types:

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

By default, both will be applied jointly, but this can be controlled by the user.

The following notation is used throughout:

$\text{Int}(x)$  denotes the integral part of  $x$

$\lceil x \rceil$  = the ceiling function, that is, the smallest integer  $n$  such that  $n - 1 \leq x \leq n$

$i$  = occurrence period

$\bar{t}$  = 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  $i$

$\lambda_i$  = 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 superimposed inflation index, representing the ratio of dollar values at calendar time  $\bar{t}$  to those at calendar time 0

$g_O(i | s)$  = occurrence period superimposed inflation 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  $i$

$s_{ir}$  = 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 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 (the epoch of the  $m$ -th partial payment is  $u_{ir} + v_{ir} + d_{ir}^{(1)} + \dots + d_{ir}^{(m)}$ )

All of these quantities from  $n_i$  onward, but except  $r$ , are realizations 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 realization is  $s_{ir}$ .

## 4. Architecture

### 4.1. Claim process

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

- 4.1.1 – Claim occurrence date;
- 4.1.2 – Claim size without allowance for inflation;
- 4.1.3 – Claim notification date;
- 4.1.4 – Claim settlement date;
- 4.1.5 – Number of partial payments;
- 4.1.6 – Sizes of partial payments without allowance for inflation;
- 4.1.7 – Distribution of payments over time;
- 4.1.8 – Claim inflation.

The modules have been listed here in the logical order of **SynthETIC** rather than temporal order. This architecture more or less follows that set out in the early papers on claim micro-models (Arjas, 1989; Jewell, 1989; Norberg, 1993, 1999; Hesselager, 1994). The following paragraphs provide full detail of the architecture of each module, though each may be re-coded by the user if desired (see Section 4.2).

The **default version** of the package encodes the modules described in the remainder of the present sub-section, in conjunction with the parameterization set out in Appendix A.

#### 4.1.1. Claim occurrence date

Occurrence period  $i$  covers calendar period  $(i - 1, i]$ . It is assumed by default that  $N_i \sim \text{Poisson}(E_i \lambda_i)$ , where  $E_i, \lambda_i$  are input parameters. The number of claims  $n_i$  is a random drawing of  $N_i$ . As indicated in Section 3,  $E_i$  is an **annual effective** parameter. In addition to the default Poisson frequency, **SynthETIC** supports a full range of user-specified alternative sampling distributions, from standard parametric ones (e.g. negative binomial) to more customized versions (e.g. non-homogeneous Poisson process); as is the case for all other modules that follow.

**SynthETIC** is also fully compatible with existing actuarial **R** packages such as **actuar** or **fitdistrplus**, which means that users can always pick a more advanced (and also more realistic) claim frequency or

severity distribution introduced by `actuar`, or simulate from a parametric model fitted on a real data set by `fitdistrplus` (Dutang et al., 2008; Delignette-Muller and Dutang, 2015). Examples are documented in the package vignette (see Section 8). The example implementation in Section 6 is specified in Appendix A to work with quarterly periods, and with  $E_i = 12000$ ,  $\lambda_i = 0.03$ . The exposure for a single occurrence quarter will therefore be 3000 (exposure-years), and the expected claim frequency  $3000 \times 0.03 = 90$ .

No trend in claim frequency over the occurrence period has been assumed, and so the occurrence time of claim  $r$  is a realization  $u_{ir}$  of a random variable  $U_{ir}$  simulated as  $U_{ir} \sim \text{Uniform}(i - 1, i]$ ,  $r = 1, 2, \dots, N_i$ .

In principle, the assumption of uniformity would conflict with fractional  $E_i$  if the position of these exposures within the occurrence period were known. This conflict has been disregarded for the purpose of the default version of `SynthETIC`, but is capable of modification if the conflict is material.

The default selections of  $E_i, \lambda_i$  are set out in Appendix A.

#### 4.1.2. Claim size without allowance for inflation

Claim size is represented as a multiple of a reference claim size, which is defined as a global parameter. This enables the simulator to switch conveniently between currencies, e.g. the reference claim size might be 1,000 USD or 1,000,000 KRW. Alternatively, some schemes of insurance define entitlements as multiples of a regulated reference claim size, and this latter may be used as the `SynthETIC` global parameter.

Claim size  $s_{ir}$  is the realization of a random variable  $S_{ir}$  with df  $F_S(s)$ , independent of  $i$  in this version of `SynthETIC`, specified on input as a function of  $s$ . This is sampled according to  $S_{ir} = F_S^{-1}(Z_{ir}^{(S)})$ , where  $Z_{ir}^{(S)} \sim \text{Uniform}(0, 1]$ . In practice, the user can choose to specify to `SynthETIC` either a distribution function (based on which `SynthETIC` applies inverse transform sampling), or directly a random generating function. The flexibility with the latter implies that the user is able to simulate from almost any model. The `SynthETIC` vignette includes an example on simulating claim sizes from a gamma GLM, and many others on alternative parametric models.

The default version of  $F_S(s)$  is set out in Appendix A.

#### 4.1.3. Claim notification date

Notification delay  $v_{ir}$  is the realization of a random variable  $V_{ir}$  with df  $F_{V|i,s}(v; i, s)$ , specified on input as a function of  $v$ , and possibly dependent on occurrence period  $i$  and claim size  $s$ . This is sampled according to  $V_{ir} = F_{V|i,s}^{-1}(Z_{ir}^{(V)})$ , where  $Z_{ir}^{(V)} \sim \text{Uniform}(0, 1]$ . Here,  $F_{V|i,s}(v; i, s)$  represents a general function form and is not limited to any particular distribution form (e.g. a Weibull distribution, as described in Appendix A). In practice, this means the user is free to choose any distribution model of their liking (with parameters possibly dependent on occurrence period  $i$  and claim size  $s$ ), as is the case for all other modules that follow (but with a slight modification on the dependence structure).

The default version of  $F_{V|i,s}(v; i, s)$  is set out in Appendix A.

#### 4.1.4. Claim settlement date

Settlement delay  $w_{ir}$  is the realization of a survival random variable  $W_{ir}$  with survival function  $\bar{F}_{W|i,s}(w; i, s)$ , specified on input as a function of  $w$ , and possibly dependent on occurrence period  $i$  and claim size  $s$ , and with survival function  $\bar{F}_{W|i,s} = 1 - F_{W|i,s}$  for df  $F_{W|i,s}$ . This is sampled according to  $W_{ir} = F_{W|i,s}^{-1}(Z_{ir}^{(W)})$ , where  $Z_{ir}^{(W)} \sim \text{Uniform}(0, 1]$ .

The default version of  $F_{W|i,s}(w; i, s)$  is set out in Appendix A.

#### 4.1.5. Number of partial payments

The number of partial payments  $m_{ir}$  is the realization of a random variable  $M_{ir}$  with df  $F_{M|s}(m; s)$ , specified on input as a function of  $m$ , and possibly dependent on claim size  $s > 0$ . This is sampled according to  $M_{ir} = F_{M|s}^{-1}(Z_{ir}^{(M)})$ , where  $Z_{ir}^{(M)} \sim \text{Uniform}(0, 1]$ .

The default version of  $F_{M|s}(m; s)$  is set out in Appendix A.

#### 4.1.6. Sizes of partial payments without allowance for inflation

Consider claims from an Auto Liability line of business, for example. Such a claim will usually consist of:

- (1) (possibly) some small payments such as police reports, medical consultations and reports;
- (2) Some more substantial payments such as hospitalization, specialist medical procedures, equipment (e.g. prosthetics);
- (3) A final settlement with the claimant;
- (4) A smaller final payment, usually covering legal costs.

The settlement with the claimant will typically account for the bulk of the claim cost, and the final payment will usually increase (not necessarily linearly) with the size of the settlement.

Claims in a number of other lines of business exhibit a similar structure, albeit with possible differences in the types of payment made. For example, in Auto Collision Damage, (1) might include tow-truck costs, and (2) might include replacement hire car costs. Payments of type (4) might be negative and account for third party recoveries and salvage.

Some claims may be trivial, involving only preliminary costs of types (1) and (2). It is assumed by default that these claims are characterized by  $m_{ir} \leq 3$ , and that the more substantial claims, with payments of types (3) and (4), are characterized by  $m_{ir} \geq 4$ .

**Case  $m_{ir} \geq 4$ .** In accordance with the above commentary, payments of types (3) and (4), i.e. the last two payments in respect of the claim, are simulated first. Initially the sum of the two is simulated, and then this amount is apportioned between the two payments.

The sum of the two payments are sampled, as a proportion of claim size  $s_{ir}$ , according to  $1 - \left[ P_{ir}^{(m_{ir}-1)} + P_{ir}^{(m_{ir})} \right] \sim F_{P|s}^{-1} \left( Z_{ir}^{(\bar{P})} \right)$ , where  $Z_{ir}^{(\bar{P})} \sim \text{Uniform}(0, 1]$ . The default version of  $F_{P|s}$  is set out in Appendix A.

Let the proportion of the total represented by the settlement be denoted

$$\bar{Q}_{ir} = \frac{P_{ir}^{(m_{ir}-1)}}{P_{ir}^{(m_{ir}-1)} + P_{ir}^{(m_{ir})}},$$

which is sampled according to  $\bar{Q}_{ir} \sim F_{\bar{Q}|s}^{-1} \left( Z_{ir}^{(\bar{Q})} \right)$ , where  $Z_{ir}^{(\bar{Q})} \sim \text{Uniform}(0, 1]$ . The default version of  $F_{\bar{Q}|s}$  is set out in Appendix A.

This leaves the proportion  $1 - \left[ P_{ir}^{(m_{ir}-1)} + P_{ir}^{(m_{ir})} \right]$  of  $s_{ir}$  to be accounted for by partial payments  $s_{ir}^{(1)}, s_{ir}^{(2)}, \dots, s_{ir}^{(m_{ir}-2)}$ .

The proportions  $p_{ir}^{(m)}, m = 1, \dots, m_{ir} - 2$  require normalization so that

$$\sum_{m=1}^{m_{ir}-2} p_{ir}^{(m)} = 1 - \left( p_{ir}^{(m_{ir}-1)} + p_{ir}^{(m_{ir})} \right).$$

They are first simulated in unnormalized form  $\hat{p}_{ir}^{(m)}$  and these quantities then normalized. Thus,  $\hat{p}_{ir}^{(m)}$  is the realization of a random variable  $\hat{P}_{ir}^{(m)} \sim F_{\hat{P}|m}^{-1} \left( Z_{ir}^{\hat{P}|m} \right)$ , where  $Z_{ir}^{\hat{P}|m} \sim \text{Uniform}(0, 1]$ . The default form of  $F_{\hat{P}|m}$  is set out in Appendix A.

Then normalization takes the form

$$p_{ir}^{(m)} = \left[ 1 - \left( p_{ir}^{(m_{ir}-1)} + p_{ir}^{(m_{ir})} \right) \right] \frac{\hat{p}_{ir}^{(m)}}{\hat{p}_{ir}^{(1)} + \dots + \hat{p}_{ir}^{(m_{ir}-2)}}, \quad m = 1, \dots, m_{ir} - 2.$$

A summary of the construction of partial payments is as follows:

- Simulate  $p_{ir}^{(m_{ir}-1)} + p_{ir}^{(m_{ir})}$ .
- Simulate  $\bar{q}_{ir}$ .
- Calculate  $p_{ir}^{(m_{ir}-1)} = \bar{q}_{ir} (p_{ir}^{(m_{ir}-1)} + p_{ir}^{(m_{ir})})$ ,  $p_{ir}^{(m_{ir})} = (1 - \bar{q}_{ir}) (p_{ir}^{(m_{ir}-1)} + p_{ir}^{(m_{ir})})$ .
- Simulate  $\hat{p}_{ir}^{(m)}$ ,  $m = 1, \dots, m_{ir} - 2$ .
- Normalize to obtain  $p_{ir}^{(m)}$ ,  $m = 1, \dots, m_{ir} - 2$ .
- Calculate  $s_{ir}^{(m)} = p_{ir}^{(m)} s_{ir}$ ,  $m = 1, \dots, m_{ir}$ .

**Case  $m_{ir} = 2$  or  $3$ .** The calculations proceed as above, except without settlement and associated partial payments. Thus,

- for actual  $m_{ir} = 2$ , proceed as if  $m_{ir} = 4$  with  $p_{ir}^{(m_{ir}-1)} + p_{ir}^{(m_{ir})} = 0$ ;
- for actual  $m_{ir} = 3$ , proceed as if  $m_{ir} = 5$  with  $p_{ir}^{(m_{ir}-1)} + p_{ir}^{(m_{ir})} = 0$ .

**Case  $m_{ir} = 1$ .** For this degenerate case,  $s_{ir}^{(1)} = s_{ir}$  of course.

The above sets up the simulation of partial payment sizes based on the number of partial payments,  $m_{ir}$ . While we believe this structure is generalizable to many lines of businesses, users of **SynthETIC** can easily amend this structure by feeding in an alternative simulation model.

#### 4.1.7. Distribution of payments over time

As noted in Section 3, the epoch of the  $m$ -th partial payment is  $u_{ir} + v_{ir} + d_{ir}^{(1)} + \dots + d_{ir}^{(m)}$ . Let this be denoted  $\bar{t}_{ir}^{(m)}$ . The payment period in which this payment falls is  $\lceil u_{ir} + v_{ir} + d_{ir}^{(1)} + \dots + d_{ir}^{(m)} \rceil$ .

The  $d_{ir}^{(m)}$  require normalization so that  $d_{ir}^{(1)} + \dots + d_{ir}^{(m_{ir})} = w_{ir}$ , and so their simulation parallels that of the  $p_{ir}^{(m)}$  just above. Thus,  $\hat{d}_{ir}^{(m)}$  is the realization of a random variable  $\hat{D}_{ir}^{(m)} \sim F_{\hat{D}|m}^{-1} \left( Z_{ir}^{\hat{D}|m} \right)$ , where  $Z_{ir}^{\hat{D}|m} \sim \text{Uniform}(0, 1]$ . The default form of  $F_{\hat{D}|m}$  is set out in Appendix A.

Normalization takes the form

$$d_{ir}^{(m)} = w_{ir} \frac{\hat{d}_{ir}^{(m)}}{\hat{d}_{ir}^{(1)} + \dots + \hat{d}_{ir}^{(m_{ir})}}, \quad m = 1, \dots, m_{ir}.$$

#### 4.1.8. Claim inflation

The actual dollar value of a constant dollar partial payment is

$$s_{ir}^{*(m)} = s_{ir}^{(m)} \frac{f \left( t_{ir}^{(m)} \right)}{f(1)} \frac{g_O(i | s_{ir})}{g_O(1 | s_{ir})} \frac{g_C \left( t_{ir}^{(m)} | s_{ir} \right)}{g_C(1 | s_{ir})}.$$

#### 4.1.9. Out-of-bounds transactions

For the purpose of the present sub-section, a transaction includes occurrence, notification, settlement, or a payment. Occasionally, simulated transactions will be out-of-bounds, i.e. take place beyond the end of development period  $I$ , where development periods are numbered  $1, 2, \dots$ .

In these cases, the simulated epoch of occurrence of the transaction is maintained throughout the simulation of details of the claim concerned, other than in the exceptions noted below. For example, if settlement occurs at development time  $j > I$ , and delays between partial payments depend on settlement delay, then the simulated value of  $j$  will be used in the simulation of inter-partial delays. Only at the stage where transactions are assigned to development periods for the purpose of either addition of inflation or (optionally) tabulation is the epoch of occurrence varied. In these circumstances, the transaction is deemed to have

occurred at the end of development period  $I$ . For tabulation, the user will also have the option to leave the out-of-bound transactions in a separate “tail” cell, from which they can decide whether to aggregate them in the maximum development period or to expand the triangle.

Accumulation into the maximum development period will cause some concentration of transactions at the end of development period  $I$ . Usually, the concentration will be small to the extent of virtual immateriality. If the effect is material, this may indicate that the user’s selection of parameters determining transaction times is not well matched to the maximum number of development periods allowed, and consideration might be given to changing one or the other.

#### 4.2. Modularity

**SynthETIC** has been structured so as to generate a portfolio of claims that loosely resemble those from an Auto Liability portfolio with which one of the authors is familiar. The latter is referred to here as the “**reference portfolio**”, and has been discussed in various earlier papers (Taylor et al., 2003; Taylor and McGuire, 2004; McGuire, 2007; McGuire et al., 2018).

The structure of the default version set out in Section 4.1 and Appendix A is quite general, and many users should be able to adopt it with changes other than in its parameters (the numerical coefficients in its algebraic structure). Slightly less simple, but still relatively easily effected, would be changes to the functional forms of sampling distributions in the default structure. Indeed, though not always encouraged, users can also add or remove intra-model dependencies (described in Sections 4.1 and 4.3) which form part of **SynthETIC**’s default architecture. Examples of such implementations are documented in the package vignette. However, we also recognise that there remain other cases in which some change of structure is required.

For this reason, **SynthETIC** has been coded in modular form. Its eight modules are listed as 4.1.1 to 4.1.8. The coding of any one is independent of all others so that the user may unplug any one and replace it with a version modified to his/her own purpose. **SynthETIC** includes functions to convert the replacement to a format that can be easily integrated with the later modules (see examples in the **SynthETIC** vignette).

It is essential to note, however, that the module sequence 4.1.1 to 4.1.8 should be maintained. The reason is that this sequence reflects assumed dependencies whereby any specific quantities generated by any specific module may be dependent on those from prior modules. Examples are given in [4.3.1] to [4.3.7].

#### 4.3. Data features

The reference portfolio is complex, containing a number of features that can cause a certain degree of awkwardness in modelling. These are as follows.

- [4.3.1] Distribution of settlement delay depends on claim size.
- [4.3.2] Distribution of settlement delay also, conditional on claim size, varies from one occurrence period to another.
- [4.3.3] Claim sizes are subject to payment period superimposed inflation.
- [4.3.4] Rates of payment period superimposed inflation vary from one payment period to another.
- [4.3.5] Rates of payment period superimposed inflation also vary with claim size.
- [4.3.6] A legislative change occurred that affected the claim experience of subsequent occurrence periods (occurrence period superimposed inflation).
- [4.3.7] The legislative change affected claims differentially according to pre-change claim size.

All of these features have been incorporated in **SynthETIC**.

## 5. Relation to prior literature

### 5.1. *Claim simulation literature*

The literature contains a few predecessor simulators, namely: CAS Loss Simulation Model Working Party (2007); Harej, Gächter, and Jamal (2017, ASTIN Working Party on Individual Claim Development with Machine Learning); and Gabrielli and Wüthrich (2018). These are discussed in the following sub-sections, where their major differences with **SynthETIC** are identified. A summary of those is provided in Table 5.1.

**SynthETIC** is the only simulator, to our knowledge, that includes all the desirable features outlined in Section 2. Furthermore, it best offers a continuous spectrum of complexity, which is crucial to assess the strengths and weaknesses of a new model; see also Section 1 for further discussion of this point.

|                                      | SynthETIC<br>(this paper)   | CAS model<br>(see Section 5.1.1)  | ASTIN WP<br>(see Section 5.1.2)  | Gabrielli-Wüthrich<br>(see Section 5.1.3)   |
|--------------------------------------|---|---|--|---|
| Availability                         | Open source: GitHub and CRAN  | Open source   | Not available  | Available online  |
| Inspiration                          | Auto Liability portfolio, but easily adjustable to most other LoBs (via parametric, rather than structural, changes)  | Auto experience in US   | Not disclosed  | Closely calibrated to four unspecified LoBs (suspected Liability) of an unspecified real insurance portfolio  |
| Time definition                      | Fully flexible in terms of time length and granularity  | Discrete dates  | Years  | Years (12)  |
| Claim size (severity) distributions  | Fully flexible, any distribution or GLM. Fully compatible with <code>actuar</code> . Covariates are allowed in very flexible way<br>See Section 4.1.2   | Select from prescribed set, options to modify to allow both zero claims and deductibles   | Lognormal with prescribed location and dispersion parameters                             | Lognormal with location as a linear function of output of the prior layer of neurons. Covariates (albeit restricted to a small prescribed set) are allowed (but their effects on the simulated data are also prescribed) and nonlinear trends can be incorporated (albeit implicit) |
| Notification and settlement event    | Notification and settlement delay distributions can be specified by user, but package includes default distributions. See Sections 4.1.3-4.1.4  | Automatic fit from prescribed set of distributions, or empirical distribution   | Not dealt with   | Notification simulated from a prescribed neural network, settlement delays sampled from a given categorical distribution  |
| Partial payments                     | Explicit random number (see Section 4.1.5), with sizes also simulated according to a realistic pattern (see Section 4.1.6)  | No provision  | Indirectly   | Deterministic payment pattern   |
| Claim development                    | Any tail length is possible (including very long)<br>See Section 4.1.7  | Calculated using user-specified year-to-year development factors (randomized through the standard deviation of the factors), but only the ultimate losses are returned        | Shape $F_P(j)$ drawn from a lognormal with prescribed location and dispersion parameters | Relatively short-tailed   |
| Claim inflation                      | All rates can be set arbitrarily:<br>- Base inflation via quarterly rates<br>- Superimposed inflation via both payment periods and occurrence periods<br>See Section 4.1.8  | Via "severity index"  | No explicit allowance  | Not explicitly, but implicitly through trend  |
| Possible model internal dependencies | Explicit modelling In the default version these include:<br>- distribution of settlement delay depends on claim size<br>- notification delay may depend on occurrence period and claim size<br>- settlement delay may depend on occurrence period and claim size<br>- number of partial payments may depend on claim size<br>- size of partial payments depend on the sizes of those preceding<br>Users can add or remove intra-model dependencies. | Via correlation only, including:<br>- Claim size and settlement delay<br>- Notification delay and settlement delay  | Frank copula between the paid and outstanding development patterns                       | - Notification delay impacts number of years with payment and claim size.<br>- Partially across development years only, as payment pattern is deterministic   |
| Extras                               | - Modular coding for enhanced flexibility<br>- Allows for structural changes such as legislative changes  | - Claim re-openings<br>- Generates distributional assumptions by fitting a prescribed set of distributions on a claim file provided by the user (following a specific format) |  | - Claim recoveries<br>- Claim re-openings   |

**Table 5.1:** Comparison of existing loss reserving simulators

### 5.1.1. CAS Loss Simulation Model Working Party

The simulator incorporates controls on the following features:

- Notification delay;
- Settlement delay;
- Claim size.

Each of these three may be specified by selecting from a prescribed set of standard distributions. There are options to modify the claim size distribution to allow for either or both of zero claims and the effect of deductibles.

The simulator is open source. It is calibrated against the Auto experience of a number of US states, although the user may over-ride this calibration.

The program includes dependencies between:

- claim size and settlement delay; and
- notification delay and settlement delay.

In both cases, dependency is expressed in terms of correlation, which differs from the present simulator. Claim inflation is included by means of a user-defined “severity index”.

No provision for partial payments is included.

### 5.1.2. ASTIN Working Party on Individual Claim Development with Machine Learning

This simulator generated samples of both paid and incurred amounts. The latter is not of concern for present purposes. Although the model of claim payments is expressed in terms of continuous time, it is implemented in discrete time (years).

Let  $j$  denote continuous development time, with origin at the commencement of the occurrence period. Cumulative claim payments to development time  $j$ , denoted  $P(j)$ , are defined as  $P(j) = UF_P(j)$ , where  $U$  is the ultimate claim size and  $F_P(j)$  is the proportion of it paid by development time  $j$ .

Claim size  $U$  is sampled from a log normal distribution with prescribed location and dispersion parameters. For each  $j = 1, 2, \dots$ ,  $F_P(j)$  is drawn from a log normal distribution also with prescribed location and dispersion parameters. In particular, the location parameter at development time  $j$  takes the form  $\alpha \exp(-((j - \tau)/\lambda)^2)$  with the quantities  $\alpha, \tau, \lambda$  used to calibrate the payment pattern to short and long tail forms.

This simulator does not deal with notification or settlement events. It does not include explicit allowance for inflation.

Since  $F_P(j)$  is a continuous function of  $j$ , the simulator envisages a claim as paid continuously by infinitesimal amounts. As pointed out above, however, it is implemented in discrete time, with payments occurring in each development year.

In this sense, allowance is made for partial payments. On the other hand, each claim is just a miniature version of aggregate payment experience, subject to random perturbation. This would not be a realistic representation of partial payments for most lines of business.

### 5.1.3. Gabrielli-Wüthrich

The simulator of Gabrielli and Wüthrich (2018) generates claim payments in discrete time (years) over 12 development years. It is calibrated against a specific set of real data from four lines of business. The lines of business are not stated, though the fact that the data include “labor sector of the injured” and “part of the body injured” suggest Casualty lines, possibly Employers Liability. The code is available on-line.

The simulator generates the following features:

- Notification delay (in development years);
- Settlement delay (in development years);

- Claim size, including the proportion of zero-cost claims;
- Number of development years with payment;
- Claim payments by development year.

Less usual features included in the simulator are:

- Covariates (accident year, accident quarter, line of business, labor sector of the injured, age of the injured, part of the body injured) which differentiate claims experience;
- Claim recoveries; and
- Claim re-openings.

The inclusion of accident year and quarter in covariates enables the model to accommodate (non-linear) time trends and seasonalities across accident periods in all model components. The case of claim size is of special interest in this context. This is included as a log normal variate with location parameter expressed as a linear function of the output of the prior layer of neurons.

Hence the logarithm of claim size incorporates a non-linear trend across accident periods. The simulator does not include explicit allowance for inflation but, according to the trend just discussed, includes implicit accident year inflation. It does not include calendar period inflation.

Dependencies between some simulated quantities are included in the model. For example, the number of development years containing payment depends on the notification delay; claim size depends on notification delay and number of development years containing payment.

Dependencies between payments of different development years are included partially by virtue of the simulation of number of development years with payment. For any fixed number of these development years, a deterministic payment pattern is assumed, whence part of the dependency between payments of different development years is excluded. The authors point out that the selection of a payment pattern could be made stochastic, but this would require careful engineering of the dependencies between development years.

The parametric structure of the simulator is fixed. The output always resembles a drawing from the data set against which the simulator is calibrated. The modeller may sometimes seek greater challenge than provided by this data set.

**SynthETIC** makes useful contributions to the already excellent simulator of Gabrielli and Wüthrich (2018). First, the latter appears relatively short-tailed; Table 5 of Gabrielli and Wüthrich (2018) indicates that more than 50% of an accident year’s cost is paid in the first development year, and roughly 90% within the first three development years. Second, the data appear to conform reasonably well with the simple chain ladder structure.

The chain ladder will often perform well in relation to data that conform with its multiplicative cross-classified parametric structure; correspondingly, it will often perform poorly in relation to data that do not so conform. Thus, purpose of proposed alternative models will often be to fill the gap left by the chain ladder, i.e. analysis of those data sets that are awkward for the chain ladder. Testing of these alternative models will then require data sets that are poorly adapted to the chain ladder. **SynthETIC** was explicitly designed to facilitate such an analysis.

Similar comments can be made about the length of payment tail of a data set. Analysis is usually most difficult in the case of long-tailed data sets, and so there is a need for the simulation of these. **SynthETIC** allows the user to easily specify arbitrarily long tails.

## 5.2. Granular model literature

One family of models likely to thirst for testing data is that of **granular (or micro-) models**. These, by their nature, often attempt to model the minutiae of individual claims, and so are closely related to **SynthETIC**, which generates this detail.

A reasonably up-to-date summary of the literature of these models is given by De Felice and Moriconi (2019). They are also discussed by Taylor (2019), where it is pointed out that the elaborateness of the granular model with its many parameters is likely to cause essential intra-model dependencies to be overlooked

or replaced by bland and unrealistic assumptions. That reference gives a couple of specific examples, one of which relates to claim payments.

Section 6.2.2 illustrates this particular issue numerically. Section 4.3 lists a number of other dependencies built into SynthETIC.

## 6. Example implementation of SynthETIC

### 6.1. Claim process components

An example simulation has been performed in accordance with the detailed specification set out in Appendix A. The generated experience covers 40 occurrence quarters, each tracked for 40 development quarters.

The principal features of the experience are similar to those of the reference portfolio, as set out in [4.3.1] to [4.3.7], though slightly less extreme than that case. All of those features are present in the simulated experience. Some specific detail of relevance is the following:

- 6.1.1. Settlement delays, in addition to depending on claim size, decline gradually by 15% over the first 20 occurrence quarters, but is stable over subsequent quarters with the exception noted in 6.1.4.
- 6.1.2. Base inflation occurs at 2% per annum.
- 6.1.3. Payment period superimposed inflation is very high (30% per annum) for the smallest claims, but zero for claims exceeding \$200,000 in dollar values of payment quarter 1. The rate of inflation varies linearly between claim sizes of zero and \$200,000.
- 6.1.4. There was a legislative change at the end of occurrence quarter 20, which, in the main, affected smaller claims in all subsequent occurrence periods. As a result, settlement delays of claims up to \$20,000 (in dollar values of payment quarter 1), already reduced by 15% (see 6.1.1.), immediately decline by a further 20%, but this effect is gradually eroded over the next 10 occurrence periods. At the same time, claims of up to \$50,000 reduce in size. The reduction is 40% for the smallest claims, nil for claims of \$50,000, and the reduction varies linearly between these claim sizes.

Some of these features cause a steady shortening of the claim payment pattern, which is illustrated in Figure 1.

### 6.2. Numerical results


#### 6.2.1. Comparison with chain ladder

The data features set out in [4.3.1] to [4.3.7] and 6.1.1. to 6.1.4. create dramatic breaches of the chain ladder assumptions under which the expectation of any occurrence/development period cell is simply the product of an occurrence period parameter and a development period parameter.

This opens the chain ladder to forecast error. This has been investigated as follows:

- (a) The part of the simulated individual claim data relating to payment quarters 1 to 40 has been aggregated into a  $40 \times 40$  payment triangle.
- (b) The chain ladder has been applied to this triangle, and a forecast of outstanding claims up to development quarter 40 obtained (in fact, there is little claim activity beyond development quarter 40).
- (c) This forecast has been compared with the “actual” amount of outstanding claims, simulated for payment quarters 41 to 79.

The results are set out in Table 6.1 where the chain ladder is seen to perform very poorly. This is perhaps unsurprising in view of the data features and the extent to which they breach chain ladder assumptions. Data sets such as this are useful for testing models that endeavour to represent data outside the scope of the chain ladder. See Section 7 for further comment.



**Figure 1:** Development pattern of SynthETIC’s test implementation

### 6.2.2. Intra-model dependencies

The general structure of payments in respect of a specific claim is described in Section 4.1. Here, larger claims are envisaged as involving a settlement payment, usually shortly before closure, and with this payment ending to dominate other payments in respect of this claim. Smaller claims do not exhibit this feature.

One effect of this is that, if a large  $m$ -th partial payment is experienced, then the likelihood that the next payment in respect of the same claim will also be large is very much reduced at larger values of  $m$ . Likewise, if only small or medium payments have been observed, then the likelihood of a larger payment in the future is increased.

This stands in contrast with some granular models, which assume independence between the sizes of partial payments. This creates difficulties for the granular model, since the incorporation of the required dependencies may be awkward without the assumption of a claim payment structure such as that used in the simulator under discussion here.

The type of dependency described above is illustrated in Table 6.2, which is an abridged summary of the simulated data set. It displays, for selected values of  $m$  and various size ranges of  $m$ -th partial payments, the average size of the  $(m + 1)$ -th partial payment, where inflation has been excluded throughout. The table confirms the expected dependencies between the sizes of successive partial payments.

**Table 6.1:** Forecast claims costs based on synthetic data set

| Occurrence quarters | Estimated loss reserve |                       |                               |
|---------------------|------------------------|-----------------------|-------------------------------|
|                     | Target (simulator)     | Chain ladder estimate | Ratio: chain ladder to target |
|                     | \$M                    | \$M                   | %                             |
| 1 to 10             | 9.2                    | 9.5                   | 3                             |
| 11 to 20            | 26.2                   | 54.6                  | 109                           |
| 21 to 25            | 26.1                   | 68.1                  | 161                           |
| 26 to 30            | 61.8                   | 103.4                 | 67                            |
| 31                  | 20.5                   | 22.4                  | 9                             |
| 32                  | 24.4                   | 34.8                  | 43                            |
| 33                  | 27.8                   | 35.5                  | 27                            |
| 34                  | 30.1                   | 44.5                  | 48                            |
| 35                  | 26.4                   | 53.2                  | 101                           |
| 36                  | 36.2                   | 47.8                  | 32                            |
| 37                  | 33.0                   | 66.4                  | 101                           |
| 38                  | 43.1                   | 44.9                  | 4                             |
| 39                  | 39.0                   | 39.5                  | 1                             |
| 40                  | 46.2                   | 32.4                  | -30                           |
| <b>Total</b>        | <b>450.0</b>           | <b>657.1</b>          | <b>46</b>                     |

## 7. Application of SynthETIC

### 7.1. Data complexity

The synthetic data set with the features described in Section 6.1 contains substantial breaches of the chain ladder assumptions, as noted in Section 6.2.1, and, in consequence, requires complex modelling. It is, however, possible to generate simple data sets using SynthETIC. Indeed, such data sets may be easily rendered compatible with the chain ladder.

To do so, one simply notes that compatibility is achieved if the expected distribution of claim payments across development periods is the same for all occurrence periods. Hence, SynthETIC will generate chain-ladder-compatible data if:

- (a) all of its components 4.1.1 to 4.1.7 are defined to be independent of occurrence period;
- (b) base inflation and calendar period superimposed inflation in 4.1.8 occur at a constant rate per period; and
- (c) occurrence period superimposed inflation in 4.1.8 must be independent of all other components, but otherwise can be arbitrary.

Condition (b) is the case because it is known that, insertion of a constant rate of calendar period superimposed inflation in a claim triangle will change the estimated distribution of claim payments over development periods, but will also introduce a compensating change in accident period parameters so as to leave the estimated loss reserve invariant (Taylor, 2000).

Occurrence period superimposed inflation is directly reflected in the chain ladder row parameters, which can be arbitrary. Hence condition (c).

The above demonstrates that SynthETIC may be used to generate very simple or very complex data sets. There is obviously a vast range of intermediate cases, and so it follows that SynthETIC may be used to generate a collection of data sets that provide a **spectrum of complexity**. Such a collection may be

**Table 6.2:** Simulated sizes of successive partial payments (excluding inflation)

| Size of the $m$ -th partial payment | Size of $(m + 1)$ -th partial payment for $m =$ |           |      |      |       |       |
|-------------------------------------|---|-----------|------|------|-------|-------|
|                                     | 1   | 2         | 3    | 5    | 7     | 9     |
|                                     | \$K   | \$K       | \$K  | \$K  | \$K   | \$K   |
| 0 to 1                              | 0.5   | 0.9       | 1.9  | 6.6  | 0.9   | 3.9   |
| 1 to 2                              | 1.6   | 3.8       | 6.0  | 8.0  | 12.7  | 20.8  |
| 4 to 5                              | 4.5   | 12.0      | 21.8 | 31.4 | 36.4  | 53.6  |
| 8 to 10                             | 8.8   | 31.5      | 48.9 | 94.3 | 133.0 | 124.8 |
| 20 to 50                            | 27.3  | 253.9     | 55.6 | 25.5 | 4.6   | 4.4   |
| 50 to 100                           | 58.7  | 1099.0    | 8.1  | 8.0  | 8.9   | 10.1  |
| over 100                            | no claims                                       | no claims | 32.2 | 31.8 | 33.3  | 37.8  |

used to present a model under test with a steadily increasing challenge. This idea is used in McGuire et al. (2018).

Different collections may be constructed with some components held constant and others subject to complex variation. In this way, one may explore the effectiveness of a contending model in the presence of different sources of complexity.

### 7.2. Claim settlement dates

The architecture of **SynthETIC** depends heavily on claim closures. Payments in respect of claims other than small tend to be concentrated toward the settlement date in the example reported in Section 6, because the settlement payment occurs on average three months ahead of closure (see Appendix A).

This means that claim closure date is a strong marker for the occurrence of a large payment. This tends to be realistic for Casualty lines that do not pay claims as annuities (such as Workers Compensation), and even for some major Property lines (such as Auto and Home).

Reserving models that give prominence to closure count data are discussed in Taylor (1986, 2000) and also in Huang et al. (2016). The sequence of papers by Taylor et al. (2003), Taylor and McGuire (2004), McGuire (2007) and McGuire et al. (2018), based on the same data set, also used this type of model.

These models rely, of course, on accuracy of the count data, and this is sometime lacking. However, Taylor and Xu (2016) applied one of these models to US data from the NAIC (Meyers and Shi, 2011), which consisted of a variety of claim count qualities. They were able to demonstrate that the use of closure counts produced tighter forecasts than obtained without it.

Despite this, the literature has largely passed over such models. Some modelling is based on both claim amounts and claim notification counts, e.g. the **double chain ladder** (Martinez-Miranda et al., 2013). Notification counts are of obvious value in enabling estimation of ultimate numbers of incurred claims, which can serve as measures of exposures over accident years.

These exposures would usually be multiplicative in the modelling of claim amounts, with expected claim payments in a triangle cell assumed proportional to number of claims incurred in the relevant accident period. This is similar to the Payments per Claim Incurred model discussed in Taylor (1986, 2000).

However, notification counts are less adapted to modelling of any departure from this proportionality. In longer tailed lines, notification and settlement may be separated by substantial numbers of years, so the contribution of notification counts is likely to be oblique at best, and usually inferior to that of closure counts (assuming the latter to be reliable). **SynthETIC**, by its construction, linking claim size and partial payments to settlement delay, provides a suitable test environment for claim models based on claim settlement counts.

The ultimate form of claim data for ML models is transactional, which requires the modelling of partial payments. Relatively few contributions to the literature confront this issue, and those that do are sometimes bedevilled by unrealistic assumptions of independence of the sort discussed in Section 6.2.2. **SynthETIC**

includes partial payments in a realistic manner, recognizing the dependencies between them, and so provides a suitable test environment for claim models at a transactional level.

### 7.3. Real data

It is mentioned in Section 4.2 that the default version of **SynthETIC** generates data generally compatible with a real Auto Bodily Injury portfolio. The actual data from that portfolio has been successfully modelled many times, including in the literature (see, for example, the citations in Section 4.2).

It is to be noted, in this connection, that those analyses have usually involved operational time in 2 per cent intervals as a covariate. Although this is not the same as modelling individual claim data, it approaches that situation.

In this sense, **SynthETIC** is seen to generate data that is essentially real, and capable of being modelled with a good deal of complexity.

## 8. SynthETIC repository

**SynthETIC** is published as an open-source R package on the Comprehensive R Archive Network (CRAN) at <https://CRAN.R-project.org/package=SynthETIC> (Avanzi et al., 2020). The package is licensed under GPL-3. The source code is fully available and users are free to copy, modify, or redistribute the program or any of its derivative versions. The re-distribution must not impose any further restrictions on the rights granted by the GPL.

**SynthETIC** has functions to sequentially simulate each of the eight modules as outlined in Section 4. The default setting assumes probability distributions of the quantities as detailed in the Appendix A (except in the case of base inflation), but the distributional assumptions can be easily modified by users to match their specific claims experiences. Users can choose to output their simulated claims in the form of a chain-ladder square by occurrence and development periods, or alternatively a structured data set at either a claim or a payment level. The data set at the payment level can then be used to visualize the claims development over time (see e.g. Figure 1). A test data set generated under the default specification is also available as part of the package. A full demonstration of **SynthETIC**'s functionalities, including examples on alternative specifications and replacement modules, can be accessed by running `RShowDoc("SynthETIC-demo", package = "SynthETIC")` in the R console after the installation of the package.

Users can install **SynthETIC** from the CRAN repository by running in R `install.packages("SynthETIC")`. A development version of the program is also available on <https://github.com/agi-lab/SynthETIC>. The GitHub repository contains, in addition to the package code,

- A PDF version of the package reference manual (also available on <https://CRAN.R-project.org/package=SynthETIC>);
- A chain-ladder analysis of the test data set (discussed in Section 6), in an Excel spreadsheet.

## 9. Conclusions

The foregoing sections describe an individual claim simulator, newly introduced to the literature. Its code is open source, and it is modular with unpluggable and re-pluggable modules for the convenience of the user. Default modules are provided, based broadly on real data, and incorporating the data features highlighted in [4.3.1] to [4.3.7].

Three existing simulators are discussed in Section 5.1. These are all useful within specific contexts, but none contains all features [4.3.1] to [4.3.7]. The Gabrielli-Wüthrich simulator (Section 5.1.3) is perhaps the most extensive in its allowance for dependencies between the various observations on a single claim, but it is calibrated to a single data set, and so its simulations always mirror that data set.

The simulator proposed here reflects a number of desirable (but complicated) data features, and is flexible in the form of model used to generate claim experience. Data features within its modules are easily adjusted,

or the modules completely replaced. This enables the generation of a collection of data sets providing a spectrum of complexity with which to challenge a proposed model (Section 7.1).

**SynthETIC** may be of especial value in testing granular models. These sometimes include unrealistic assumptions of independence between different variates within the model. **SynthETIC**, on the other hand, contains built-in dependencies, e.g. between different partial payments.

## Acknowledgements

The authors are very grateful for comments from two referees, which led to significant improvements of the paper and **SynthETIC**, the R package that we developed for this work.

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

- Arjas, E., 1989. The claims reserving problem in non-life insurance: Some structural ideas. *ASTIN Bulletin* 19 (2), 139–152.
- Avanzi, B., Taylor, G., Wang, M., Wong, B., 2020. **SynthETIC**: Synthetic experience tracking insurance claims. <https://CRAN.R-project.org/package=SynthETIC>.
- Berquist, J. R., Sherman, R. E., 1977. Loss reserve adequacy testing: A comprehensive, systematic approach. In: *Proceedings of the Casualty Actuarial Society*. Vol. 64. pp. 123–184.
- Blier-Wong, C., Cossette, H., Lamontagne, L., Marceau, E., 2021. Machine learning in p&c insurance: A review for pricing and reserving. *Risks* 9 (1).  
URL <https://www.mdpi.com/2227-9091/9/1/4>
- CAS Loss Simulation Model Working Party, 2007. Parameterizing the loss simulation model. <https://www.casact.org/research/lsmwp/bsupaper.pdf>.
- Clemente, G. P., Spedicato, G., Schewe, F., 2014. The use of gamlss in assessing the distribution of unpaid claims reserves. *Casualty Actuarial Society E-Forum* Fall 2014.
- De Felice, M., Moriconi, F., 2019. Claim watching and individual claims reserving using classification and regression trees. *Risks* 7 (4), 1–36.
- Delignette-Muller, M. L., Dutang, C., 2015. **fitdistrplus**: An R package for fitting distributions. *Journal of Statistical Software* 64 (4), 1–34.  
URL <https://www.jstatsoft.org/v64/i04/>
- Delong, L., Lindholm, M., Wüthrich, M. V., 2020. Collective reserving using individual claims data. Available at SSRN.
- Dutang, C., Goulet, V., Pigeon, M., 2008. **actuar**: An R package for actuarial science. *Journal of Statistical Software* 25 (7), 38.  
URL <https://www.jstatsoft.org/v25/i07>
- England, P. D., Verrall, R. J., 2001. A flexible framework for stochastic claims reserving. In: *Proceedings of the Casualty Actuarial Society*. Vol. 88. pp. 1–38.
- Fisher, W. H., Lange, J. T., 1973. Loss reserve testing: a report year approach. *Proceedings of the Casualty Actuarial Society* 60, 189–207.
- Gabrielli, A., 2020. A neural network boosted double overdispersed poisson claims reserving model. *ASTIN Bulletin: The Journal of the IAA* 50 (1), 25–60.
- Gabrielli, A., Wüthrich, M. V., 2018. Individual claims history simulation machine. *Risks* 6 (2), 29.
- Gesmann, M., Murphy, D., Zhang, Y. W., Carrato, A., Wüthrich, M., Concina, F., Dal Moro, E., 2020. **ChainLadder**: Statistical methods and models for claims reserving in general insurance. R package version 0.2.11.  
URL <https://CRAN.R-project.org/package=ChainLadder>
- Harej, B., Gächter, R., Jamal, S., 2017. Individual claim development with machine learning. [http://www.actuaries.org/ASTIN/Documents/ASTIN\\_ICDML\\_WP\\_Report\\_final.pdf](http://www.actuaries.org/ASTIN/Documents/ASTIN_ICDML_WP_Report_final.pdf).
- Hesselager, O., 1994. A markov model for loss reserving. *ASTIN Bulletin* 24, 183–93.
- Huang, J., Wu, X., Zhou, X., 2016. Asymptotic behaviors of stochastic reserving: aggregate versus individual models. *European Journal of Operational Research* 249, 657–666.
- Jewell, W. S., 1989. Predicting ibnr events and delays. *ASTIN Bulletin* 19 (1).
- Kuo, K., 2019. Deeptriangle: A deep learning approach to loss reserving. *Risks* 7 (3), 97.
- Martinez-Miranda, M. D., Nielsen, J. P., Verrall, R. J., 2013. Double chain ladder and bornhuetter-ferguson. *North American Actuarial Journal* 17 (2), 101–113.
- 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.

- Meyers, G., 2015. Stochastic Loss Reserving Using Bayesian MCMC Models. Vol. 1 of CAS Monograph Series. Arlington, USA: Casualty Actuarial Society.
- Meyers, G., Shi, P., September 2011. Loss Reserving Data Pulled From NAIC Schedule P. [http://www.casact.org/research/index.cfm?fa=loss\\_reserves\\_data](http://www.casact.org/research/index.cfm?fa=loss_reserves_data).
- Mulquiney, P., 2006. Artificial neural networks in insurance loss reserving. In: Proceedings of the 9th Joint International Conference on Information Sciences (JCIS-06). Atlantis Press.  
URL <https://doi.org/10.2991/jcis.2006.67>
- Norberg, R., 1993. Prediction of Outstanding Liabilities in Non-Life Insurance. ASTIN Bulletin 23 (1), 95–115.
- Norberg, R., 1999. Prediction of Outstanding Liabilities - II Model Variations and Extensions. ASTIN Bulletin 29 (1), 5–27.
- Reid, D. H., 1978. Claim reserves in general insurance. Journal of the Institute of Actuaries 105, 211–296.
- Richman, R., 2020a. Ai in actuarial science – a review of recent advances – part 1. Annals of Actuarial Science, 1–23.
- Richman, R., 2020b. Ai in actuarial science – a review of recent advances – part 2. Annals of Actuarial Science, 1–29.
- Taylor, G., 2000. Loss Reserving: An Actuarial Perspective. Huebner International Series on Risk, Insurance and Economic Security. Kluwer Academic Publishers.
- 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. In: Spring 2004 Meeting of the Casualty Actuarial Society, Colorado Springs, Colorado.
- Taylor, G., McGuire, G., Greenfield, A., 2003. Loss reserving: Past, present and future, university of Melbourne Research Paper.
- Taylor, G. C., 1986. Claims reserving in non-life insurance. North-Holland, Amsterdam, Netherlands.
- Taylor, G. C., Xu, J., 2016. An empirical investigation of the value of finalisation count information to loss reserving. Variance in press.
- Wüthrich, M. V., 2018a. Machine learning in individual claims reserving. Scandinavian Actuarial Journal 2018 (6), 465–480.
- Wüthrich, M. V., 2018b. Neural networks applied to chain-ladder reserving. European Actuarial Journal 8 (2), 407–436.
- Wüthrich, M. V., Merz, M., 2008. Stochastic claims reserving methods in insurance. John Wiley & Sons, Chichester.

## A. Parametrizations

The following table displays the formal parameterization of modules 4.1.1 to 4.1.8 for the example of Section 6.

| Parameter type                                | Functional form   | Numerical parameters  |
|---|---|---|
| <b>Global</b>                                 |   | Time unit = 1/4 (calendar quarter)<br>Reference claim size = 200,000 <sup>1</sup>   |
| <b>Claim occurrence</b> <sup>2</sup>          |   | $I = 40$<br>$E_i = 12000$<br>$\lambda_i = 0.03$   |
| <b>Claim size</b> <sup>1</sup>                | Power-normal  | $S_{ir}^{0,2} \sim N(9.5, 3)$   |
| <b>Claim notification</b> <sup>1</sup>        | Weibull   | Mean = $\min(3, \max(1, 2 - \frac{1}{3} \ln(s_{ir}/100000)))$<br>Coefficient of variation = 70%   |
| <b>Claim closure</b> <sup>1</sup>             | Weibull   | Mean = $a(i) \min(25, \max(1, 6 + 4 \ln(s_{ir}/20000)))$<br>Coefficient of variation = 60%<br>where $a(i) = \max(0.85, 1 - 0.0075i)$ , subject to the over-riding condition that, for $s_{ir} < 20000$ and $i \geq 21$ , $a(i) = \min(0.85, 0.65 + 0.02(i - 21))$   |
| <b>Partial payments: Number</b> <sup>1</sup>  |   | For $s_{ir} \leq 7500$ , $\text{Prob}(M_{ir} = 1) = \text{Prob}(M_{ir} = 2) = 1/2$<br>For $7500 \leq s_{ir} \leq 15000$ , $\text{Prob}(M_{ir} = 2) = 1/3$ , $\text{Prob}(M_{ir} = 3) = 2/3$<br>For $s_{ir} > 15000$ , distribution of $M_{ir}$ is geometric, with minimum 4 and mean = $\min(8, 4 + \ln(s_{ir}/15000))$ |
| <b>Partial payments: Amounts</b> <sup>1</sup> | $F_{P s} = \text{Beta}$<br>$F_{\bar{Q} s} = \text{Beta}$<br>$F_{\hat{P} m} = \text{Beta}$ | Mean = $1 - \min(0.95, 0.75 + 0.04 \ln(s_{ir}/20000))$<br>Coefficient of variation = 20%<br>Mean = 0.9<br>Coefficient of variation = 3%<br>Mean = $(1 - (p_{ir}^{(m_{ir}-1)} + p_{ir}^{(m_{ir})})) / (m_{ir} - 2)$<br>Coefficient of variation = 10%  |
| <b>Distribution of payments over time</b>     | $F_{\hat{D} m} = \text{Weibull}$  | For $m_{ir} \geq 4$ and $m = m_{ir}$ ,<br>Mean = 3 months (converted to the relevant time units)<br>Coefficient of variation = 20%<br>For $m_{ir} \geq 4$ and $m < m_{ir}$ , or $m_{ir} < 4$ ,<br>Mean = mean claim closure delay (from notification) / $m_{ir}$<br>Coefficient of variation = 35%                      |
| <b>Base inflation</b>                         | $f(\bar{t})$  | = $(1 + \alpha)^{\bar{t}}$ , where $\alpha$ is equivalent to an increase of 2% p.a., allowing for the relevant time units   |
| <b>Superimposed inflation</b> <sup>1</sup>    | $g_O(u s)$<br>$g_C(\bar{t} s)$  | = 1 for $u \leq 20 = 1 - 0.4 \max(0, 1 - s/50000)$ for $u > 20$<br>= $(1 + \beta(s))^{\bar{t}}$ with $\beta = \gamma \max(0, 1 - s/200000)$ and $\gamma$ is equivalent to a 30% p.a. inflation rate, allowing for the relevant time units   |

Notes:

1. This component is defined in terms of claim size. The definition here displays claim size in raw (uninflated) units. The inputs to the example application of SynthETIC, on the other hand, express claim sizes as multiples of a reference claim size equal to 200,000. For example, the amount of 100,000 that appears in the definition of claim notification delay is expressed in SynthETIC as  $0.5 \times 200,000$ .
2. Exposure and claim frequency are annual effective parameters.

## B. Claim development triangles of the example implementation

For space considerations, below we show only the claim development triangles on a **yearly** basis; however, the underlying data is calculated based on quarterly development pattern and is available on the **SynThETIC** repository (see Section 8).

| Accident Year | Development Year |               |               |               |               |               |               |               |               |               |
|---------------|------------------|---------------|---------------|---------------|---------------|---------------|---------------|---------------|---------------|---------------|
|               | 1                | 2             | 3             | 4             | 5             | 6             | 7             | 8             | 9             | 10            |
| 1             | 1,066            | 9,452         | 18,927        | 30,287        | 42,525        | 56,468        | 62,876        | 69,553        | 73,852        | <b>76,152</b> |
| 2             | 1,463            | 9,836         | 25,669        | 37,002        | 46,925        | 57,862        | 64,484        | 67,626        | <b>74,977</b> | 78,175        |
| 3             | 3,021            | 15,491        | 34,694        | 49,257        | 59,293        | 66,771        | 72,653        | <b>78,445</b> | 81,514        | 82,516        |
| 4             | 2,247            | 14,847        | 33,546        | 50,969        | 60,592        | 74,500        | <b>79,288</b> | 81,996        | 83,271        | 85,974        |
| 5             | 1,892            | 20,421        | 41,594        | 62,943        | 77,215        | <b>87,601</b> | 92,755        | 96,667        | 97,922        | 101,723       |
| 6             | 4,335            | 22,748        | 47,748        | 67,447        | <b>84,501</b> | 94,869        | 100,668       | 104,826       | 106,442       | 107,307       |
| 7             | 2,015            | 21,542        | 45,726        | <b>67,976</b> | 81,616        | 91,775        | 99,333        | 104,418       | 106,847       | 106,901       |
| 8             | 4,389            | 27,606        | <b>55,346</b> | 83,075        | 99,509        | 109,076       | 115,446       | 119,931       | 122,561       | 124,959       |
| 9             | 2,955            | <b>32,372</b> | 63,773        | 92,863        | 114,928       | 131,745       | 138,791       | 145,341       | 148,514       | 152,235       |
| 10            | <b>4,896</b>     | 43,386        | 78,378        | 113,860       | 133,274       | 150,229       | 162,189       | 164,782       | 165,262       | 166,216       |

**Table A:** Cumulative claim development triangle of simulated claims (\$000)

| Accident Year | Development Year |               |               |               |               |               |               |               |               |               |
|---------------|------------------|---------------|---------------|---------------|---------------|---------------|---------------|---------------|---------------|---------------|
|               | 1                | 2             | 3             | 4             | 5             | 6             | 7             | 8             | 9             | 10            |
| 1             | 1,066            | 9,452         | 18,927        | 30,287        | 42,525        | 56,468        | 62,876        | 69,553        | 73,852        | <b>76,152</b> |
| 2             | 1,463            | 9,836         | 25,669        | 37,002        | 46,925        | 57,862        | 64,484        | 67,626        | <b>74,977</b> | 78,512        |
| 3             | 3,021            | 15,491        | 34,694        | 49,257        | 59,293        | 66,771        | 72,653        | <b>78,445</b> | 84,408        | 88,368        |
| 4             | 2,247            | 14,847        | 33,546        | 50,969        | 60,592        | 74,500        | <b>79,288</b> | 85,157        | 91,730        | 96,196        |
| 5             | 1,892            | 20,421        | 41,594        | 62,943        | 77,215        | <b>87,601</b> | 96,900        | 104,228       | 112,172       | 117,358       |
| 6             | 4,335            | 22,748        | 47,748        | 67,447        | <b>84,501</b> | 100,938       | 111,662       | 119,985       | 129,126       | 135,184       |
| 7             | 2,015            | 21,542        | 45,726        | <b>67,976</b> | 84,535        | 100,969       | 111,681       | 120,014       | 129,151       | 135,200       |
| 8             | 4,389            | 27,606        | <b>55,346</b> | 81,449        | 101,431       | 121,211       | 134,165       | 144,180       | 155,198       | 162,506       |
| 9             | 2,955            | <b>32,372</b> | 71,595        | 105,756       | 131,734       | 157,444       | 174,346       | 187,294       | 201,656       | 211,328       |
| 10            | <b>4,896</b>     | 31,943        | 66,096        | 95,894        | 118,288       | 140,843       | 155,004       | 166,841       | 179,200       | 186,893       |

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