risks

MDPI

Article

Evaluating Transition Rules for Enhancing Fairness in Bonus–Malus Systems: An Application to the Saudi Arabian Auto Insurance Market

Asrar Alyafie [1,2], Corina Constantinescu [1] and Jorge Yslas [1,*]

[1] Department of Mathematical Sciences, Institute for Financial and Actuarial Mathematics, University of Liverpool, Liverpool L69 7ZL, UK; a.a.alyafie@liverpool.ac.uk (A.A.); c.constantinescu@liverpool.ac.uk (C.C.)
[2] Department of Mathematics and Statistics, College of Science, University of Jeddah, Jeddah 21959, Saudi Arabia
* Correspondence: jorge.yslas-altamirano@liverpool.ac.uk

Abstract: A Bonus–Malus System (BMS) is a ratemaking mechanism used in insurance to adjust premiums based on a policyholder's claim history, with the goal of segmenting risk profiles more accurately. A BMS typically comprises three key components: the number of BMS levels, the transition rules dictating the movements of policyholders within the system, and the relativities used to determine premium adjustments. This paper explores the impact of modifications to these three elements on risk classification, assessed through the mean squared error. The model parameters are calibrated with real-world data from the Saudi auto insurance market. We begin the analysis by focusing on transition rules based solely on claim frequency, a framework in which most implemented BMSs work, including the current Saudi BMS. We then consider transition rules that depend on frequency and severity, in which higher penalties are given for large claim sizes. The results show that increasing the number of levels typically improves risk segmentation but requires balancing practical implementation constraints and that the adequate selection of the penalties is critical to enhancing fairness. Moreover, the study reveals that incorporating a severity-based penalty enhances risk differentiation, especially when there is a dependence between the claim frequency and severity.

Keywords: Bonus–Malus system; optimal relativities; transition rules; random effects

check for updates

Received: 18 December 2024
Revised: 5 January 2025
Accepted: 14 January 2025
Published: 20 January 2025

Citation: Alyafie, Asrar, Corina Constantinescu, and Jorge Yslas. 2025. Evaluating Transition Rules for Enhancing Fairness in Bonus–Malus Systems: An Application to the Saudi Arabian Auto Insurance Market. *Risks* 13: 18. https://doi.org/10.3390/risks13010018

Copyright: © 2025 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (https://creativecommons.org/licenses/by/4.0/).

1. Introduction

A Bonus–Malus System (BMS) is an experience rating mechanism in auto insurance that adjusts premiums based on a policyholder's claim history, rewarding claim-free periods with discounts (bonuses) and penalising claims with surcharges (maluses). In practice, a BMS consists of a finite number of levels, each associated with a specific relative premium. New policyholders typically begin at a set initial level. At the end of each policy year, their level is revised according to predefined transition rules, which depend on their claim history during that year. The base premium is then adjusted proportionally according to the corresponding relativity of the current level to determine the premium charged. While the base premium is calculated based on the policyholder's observable characteristics, it is impossible to account for all risk-relevant factors. Thus, the BMS relativities work as adjustments to address this residual heterogeneity, ensuring that the final charged premiums better reflect the actual risk posed by each policyholder. In this context, one of the fundamental aspects of designing a BMS is determining the premium relativities. According

to Norberg (1976), given a fixed number of levels, starting level, and transition rules, the optimal relativity for a given level is obtained by maximising the asymptotic predictive accuracy. More specifically, the optimal relativities minimise the mean squared error (MSE) between an (infinitely old) policy's expected aggregated claim and its premium. It is worth mentioning that although the use of a quadratic loss function has been the traditional way of computing premium relativities, other loss functions have been considered in the literature. For example, the exponential loss function has been considered in Denuit and Dhaene (2001), while the absolute value loss function was studied in Heras et al. (2002, 2004). Furthermore, other alternatives, such as the use of ruin probabilities, have been proposed as a measure to evaluate premium relativities and even transition rules (see, e.g., Afonso et al. 2017; Asmussen et al. 2021; Li et al. 2015). In addition to aligning premiums more closely with individual risk profiles, BMSs have also proven to incentivise safer driving behaviour (see, e.g., Dionne and Ghali 2005) and alleviate insurance fraud (cf. Moreno et al. 2006), emphasising their relevance in the insurance sector. We refer to, e.g., Denuit et al. (2007); Lemaire (1995) for comprehensive accounts of BMSs.

Despite their advantages, traditional BMS designs face notable limitations. As observed in Lemaire (1995), most implemented BMSs consider only the claim frequency to determine the movement of policyholders in the different risk levels. In fact, the author only identified the Korean BMS as one of the instances where claim sizes are considered. This means that in the vast majority of BMSs, claims are penalised equally, regardless of their severity. In this sense, these systems are unfair to policyholders with minor claims. For example, existing research suggests that women, while potentially involved in a higher frequency of accidents, are typically associated with less severe claims compared to men (see, e.g., Al-Garawi et al. 2021). Consequently, systems that only account for claim frequency may over-penalise female drivers. The need to incorporate severities into BMS design has been identified for some time in the literature, and some solutions have been provided. Notable contributions in this direction include Frangos and Vrontos (2001), Gómez-Déniz (2016), Ni et al. (2014), and Tzougas et al. (2014), which propose methodologies to integrate severity considerations into BMS frameworks.

Additionally, conventional BMS design has typically been based on the assumption that claim frequency and claim severity are independent factors, a simplification that has been challenged by recent research. Studies such as Frees et al. (2016) and Garrido et al. (2016) provide empirical evidence of significant dependence between these two components. Despite this observation, there is still limited research regarding BM design under the assumption of frequency–severity dependence. One notable contribution in this direction is Oh et al. (2020), where the authors used a bivariate random effect to capture the dependence between claim frequency and severity. More specifically, they derived optimal relativities depending solely on the claim frequency that implicitly incorporate the impact of claim sizes through the modelled association between frequency and severity. This work was more recently extended in Oh et al. (2022) to include transition rules based on both the number and size of claims. Additionally, in Ahn et al. (2022), the frequency–severity dependence model in Oh et al. (2020) was adapted to BMSs with long memory transition rules introduced in Lemaire (1995); Pitrebois et al. (2003a).

The objective of this paper is to assess the impact of various components of a BMS on risk classification. More specifically, we examine alterations to the number of BMS levels, the transition rules, and the structure of premium relativities, measuring their effects in terms of the MSE. In addition to evaluating changes within the classical frequency-dependent framework, we extend the analysis to incorporate claim severities into the transition rules. Furthermore, we explore the influence of the dependence between claim frequency and severity on the system's performance. The motivation for this work comes

from the current developments in the Saudi Arabian auto insurance market. Firstly, this market seemed to be sudden incorporation of a group of novice female drivers without claim history, resulting from a 2017 decree by King Salman that granted the right to drive to women in Saudi Arabia for the first time. Secondly, in 2017, the Saudi Arabian Monetary Agency (SAMA) introduced a BMS for the Saudi auto insurance market with rules that depend only on the frequency. The readiness of this system to accommodate the new cohort of novice drivers was recently evaluated in Alyafie et al. (2023). Employing a range of actuarial evaluation tools, the study concluded that the Saudi BMS converges slowly towards its stationary distribution, severely penalises newly insured people, and does not adapt well to changes in claim frequency. This work highlights the need for a new BMS for the Saudi insurance market that is fairer for novice drivers. Thus, the present work aims to provide actionable insights for designing a fairer and more effective BMS for this market.

The remainder of the paper is organised as follows. Section 2 presents the general mathematical framework for describing a BMS and some specific distributional assumptions we consider. In Section 3, we study the impact of modifications to BMS characteristics on risk segmentation, focusing on systems where the transition rules are based solely on claim frequency. Section 4 extends this analysis to frequency–severity transition rules, exploring two distinct dependence structures between claim frequency and severity: independence (Section 4.1.1) and negative correlation (Section 4.1.2). Finally, Section 5 summarises our findings.

2. Setup

We consider a portfolio of short-term auto insurance contracts, where policyholders can renew or terminate their policies at the end of each policy year. Additionally, the insurer can adjust a policyholder's premium at the start of each new year based on their claim history. We let N_t be the number of claims in the tth policy year for a randomly selected policyholder in this portfolio and $(Y_{t,k})_{k\in\mathbb{N}}$ be a sequence of independent and identically distributed (iid) random variables representing the corresponding individual claim sizes. Then, we define the policyholder's aggregate claim amount in the tth policy year as

$$S_t = \begin{cases} \sum_{k=1}^{N_t} Y_{t,k}, & N_t > 0, \\ 0, & N_t = 0. \end{cases}$$

Ratemaking is the process by which insurers determine the premiums to be charged for insurance policies. It aims to assess the risk associated with each policyholder and design a tariff structure that fairly distributes the claims burden among policyholders. Traditionally, the ratemaking process involves two distinct steps. In the first step, a priori rating, insurers use observable risk classification variables, such as characteristics of the driver and the automobile, to group a portfolio of policyholders into homogeneous tariff classes. However, these a priori variables cannot fully account for all risk characteristics of the insured peopled. Nevertheless, it is reasonable to believe that a policyholder's claim data can infer some of the unobservable risk characteristics. Thus, a second step, the a posteriori rating, typically in the form of a Bonus–Malus system (BMS), is used to adjust the premiums based on individual claim experiences.

Let us consider a homogeneous group of policyholders according to their observable characteristics. To account for the residual heterogeneity, we consider a random vector $(\Theta^{[1]}, \Theta^{[2]})$, where $\Theta^{[1]}$ is a random factor affecting the claim frequency and $\Theta^{[2]}$ affects the claim severity. More specifically, we make the following marginal distributional assumptions.

For the frequency, we assume that

$$N_t \mid \Theta^{[1]} = \theta^{[1]} \sim \text{Poisson}(\lambda\theta^{[1]}),$$

where $\lambda > 0$ represents the mean frequency of the group and

$$\Theta^{[1]} \sim \text{Gamma}(\alpha, \alpha),$$

with $\alpha > 0$, so that $\mathbb{E}[\Theta^{[1]}] = 1$ and $\text{Var}[\Theta^{[1]}] = \alpha^{-1}$. Under these assumptions, we have that the marginal unconditional distribution of N_t is a Negative Binomial with a probability mass function

$$f_{N_t}(n) = \frac{\Gamma(n+\alpha)}{\Gamma(\alpha)n!}\left(\frac{\alpha}{\lambda+\alpha}\right)^{\alpha}\left(1-\frac{\alpha}{\lambda+\alpha}\right)^{n}, \quad n = 0, 1, 2, \ldots,$$

a mean $\mathbb{E}[N_t] = \mathbb{E}[\lambda\Theta^{[1]}] = \lambda$, and a variance $\text{Var}[N_t] = \lambda(1 + \alpha^{-1}\lambda)$.

Concerning the claim severity, we assume that

$$Y_{t,k} \mid \Theta^{[2]} = \theta^{[2]} \sim \text{Gamma}(\theta^{[2]}, 1),$$

and

$$\Theta^{[2]} \sim \text{Lognormal}\left(\frac{-\sigma^2}{2}, \sigma^2\right),$$

where $\sigma > 0$. Thus, $\mathbb{E}[\Theta^{[2]}] = 1$ and $\text{Var}[\Theta^{[2]}] = \exp(\sigma^2) - 1$. We note that under the above assumption, there is no explicit representation of the marginal unconditional distribution of the claim severity $Y_{t,k}$. Nevertheless, we have that $\mathbb{E}[Y_{t,k}] = \mathbb{E}[\Theta^{[2]}] = 1$ and $\text{Var}[Y_{t,k}] = \exp(\sigma^2)$. Finally, we assume that N_t and $(Y_{t,k})_{k\in\mathbb{N}}$ are conditionally independent given $(\Theta^{[1]}, \Theta^{[2]})$. We will consider different assumptions on the dependence structure of $(\Theta^{[1]}, \Theta^{[2]})$ in later sections.

Remark 1. *The above distributional assumptions are made to illustrate the methodologies described in what follows, and they can easily be changed to reflect the data at hand. Additionally, the extension to incorporate the observable characteristics into the analysis is straightforward by following, for example, the work of Pitrebois et al. (2003b). Nevertheless, to focus on the main ideas of the methods, we omit such a case.*

Every BMS consists of three key components: the number of BMS levels, including a pre-defined initial level; the transition rules that dictate how policyholders move between the levels according to their claim history; and a set of premium relativities, which are multiplicative factors applied to the base premiums to determine the final premiums. In this paper, we denote by $1, \ldots, z$ the BMS levels and by r_l, $l = 1, \ldots, z$, the corresponding premium relativities, which typically follow a non-decreasing pattern, that is, $r_1 \leq r_2 \leq \cdots \leq r_z$. Thus, level 1 corresponds to the highest bonus or discount level (lower premium), while level z represents the highest malus or surcharge level (higher premium). Additionally, we let L_t be the current level of a policyholder in the tth policy year. Then, the level for the $(t+1)$-th policy year, L_{t+1}, will be determined as a function of the current level L_t and the claim history in the tth policy year N_t, $(Y_{t,k})_{k\in\mathbb{N}}$.

Given that the next BMS level of a policyholder is determined solely by their information in the current year, the mechanism of a BMS can be modelled using a Markov chain. Specifically, let us denote by $p_{ij}(\lambda\theta^{[1]}, \theta^{[2]})$ the one-step transition probability of moving from level i to level j for a policyholder with an expected claim frequency $\lambda\theta^{[1]}$ and an expected claim severity $\theta^{[2]}$, and by $\boldsymbol{P}(\lambda\theta^{[1]}, \theta^{[2]}) = \{p_{ij}(\lambda\theta^{[1]}, \theta^{[2]})\}_{i,j=1,\ldots,z}$ the

corresponding one-step transition matrix. Then, the associated stationary distribution $\boldsymbol{\pi}(\lambda\theta^{[1]}, \theta^{[2]}) = (\pi_1(\lambda\theta^{[1]}, \theta^{[2]}), \ldots, \pi_z(\lambda\theta^{[1]}, \theta^{[2]}))$ describing the policyholder's behaviour in the long run within the system, can be obtained by solving

$$\boldsymbol{\pi}(\lambda\theta^{[1]}, \theta^{[2]}) = \boldsymbol{\pi}(\lambda\theta^{[1]}, \theta^{[2]})\boldsymbol{P}(\lambda\theta^{[1]}, \theta^{[2]}). \tag{1}$$

Building on the above, let L be the random variable representing the BM level occupied by a randomly selected policyholder in the steady-state. The distribution of L, which represents the long-term distribution of policyholders across the BMS levels, can be expressed as follows:

$$\mathbb{P}(L = l) = \int_0^\infty \int_0^\infty \pi_l(\lambda\theta^{[1]}, \theta^{[2]}) f_{\Theta^{[1]},\Theta^{[2]}}(\theta^{[1]}, \theta^{[2]}) d\theta^{[1]} d\theta^{[2]}, \quad l = 1, \ldots, z, \tag{2}$$

where $f_{\Theta^{[1]},\Theta^{[2]}}$ represents the joint density function of the random factors affecting claim frequency and severity.

Typically, the a priori rate-making process is carried out under the assumption of independence of the claim frequency and severity. Under our distributional assumptions, this translates to having a base premium of λ for all policyholders in the group. Thus, a policyholder at the BM level l pays a total premium of λr_l. Nevertheless, if we take into account the residual heterogeneity, the premium for a policyholder with expected claim frequency $\lambda\theta^{[1]}$ and expected claim severity $\theta^{[2]}$ should be set to $\lambda\theta^{[1]}\theta^{[2]}$. This motivates us to concentrate our study on the expected difference between the charged premium and the premium that should be charged (after the steady state has been reached). More specifically, we quantify this difference in terms of the mean square error (MSE)

$$\mathbb{E}\left[\left(\lambda\Theta^{[1]}\Theta^{[2]} - \lambda r_L\right)^2\right], \tag{3}$$

which can be thought of as a measure of the predictive accuracy of the system. As such, our focus here is on minimising this quantity by modifying three key features of the BMS: the number of levels, the transition rules, and the structure of the relativities. In other words, our objective is to assess how these transition rule modifications improve policyholders' classification according to their unobservable risk characteristics. Lastly, we note that minimising (3) is equivalent to minimising

$$\mathbb{E}\left[\left(\Theta^{[1]}\Theta^{[2]} - r_L\right)^2\right], \tag{4}$$

where this last expression is the one that we will utilise in this study.

3. Transition Rules Based on Frequency

In this section, we focus on BMS transition rules that depend solely on claim frequency. It is worth noting that the vast majority of implemented BMSs operate under this framework. Furthermore, these systems are typically designed under the assumption that claim frequency and severity are independent. We adopt the same assumption here, which in our model specification translates to considering $\Theta^{[1]} \perp\!\!\!\perp \Theta^{[2]}$.

Here, we will focus on the so-called $-1/+h$ transition rules, where $h \geq 1$ represents the penalty of the system. Under this mechanism, a policyholder's level decreases by one after a claim-free year, reflecting a reward for no claims. Conversely, if one or more claims

are made in a given year, the level increases by h, imposing a penalty. Mathematically, this can be expressed as

$$L_{t+1} = \begin{cases} \max\{L_t - 1, 1\}, & N_t = 0, \\ \min\{L_t + h, z\}, & N_t > 0. \end{cases}$$

To illustrate how the $-1/+h$ transition rules work, consider a system with $z = 10$ levels and $h = 2$, that is, a $-1/+2$ system. Suppose a policyholder is currently at level $L_t = 5$. If the policyholder has no claims during the present policy year, i.e., $N_t = 0$, then the level for the next year would be $L_{t+1} = 4$. On the other hand, if the policyholder has at least one claim, that is, $N_t > 0$, then they move to level $L_{t+1} = 7$ for the subsequent policy year.

We focus on the $-1/+h$ assumption not only because of its prevalence in existing BMSs but also in accordance with the later application considered. Note that under this particular instance of BM rules, the one-step transition probabilities are non-zero only for the following elements

$$p_{i,\max(i-1,1)}(\lambda\theta^{[1]}) = \exp(-\lambda\theta^{[1]}), \quad p_{i,\min(i+h,z)}(\lambda\theta^{[1]}) = 1 - \exp(-\lambda\theta^{[1]}), \quad i = 1, \ldots, z.$$

Observe that we have omitted the dependence on $\theta^{[2]}$ in the notation above, as the transition probabilities are not affected by the claim severities. Similarly, we rewrite the stationary distribution (1) as

$$\boldsymbol{\pi}(\lambda\theta^{[1]}) = (\pi_1(\lambda\theta^{[1]}), \ldots, \pi_z(\lambda\theta^{[1]})).$$

With this notation, we have that (2) takes the form

$$\mathbb{P}(L = l) = \int_0^\infty \pi_l(\lambda\theta^{[1]}) f_{\Theta^{[1]}}(\theta^{[1]}) d\theta^{[1]}, \quad l = 1, \ldots, z.$$

Finally, note that under these assumptions, (4) further simplies to

$$\mathbb{E}\left[\left(\Theta^{[1]}\Theta^{[2]} - r_L\right)^2\right] = \left(\mathbb{E}[\Theta^{[1]}]\right)^2 \mathrm{Var}[\Theta^{[2]}] + \mathbb{E}\left[\left(\Theta^{[1]} - r_L\right)^2\right].$$

In particular, the equation above implies that in this setting, minimising (4) is equivalent to minimising

$$\mathbb{E}\left[\left(\Theta^{[1]} - r_L\right)^2\right]. \tag{5}$$

3.1. Bayesian Relativities

According to Norberg (1976), once the number of classes and the transition rules are established, the optimal (or Bayesian) relativities $\mathbf{r} = (r_1, \ldots, r_z)$, should be determined by maximising the asymptotic predictive accuracy. In other words, the relativities should be chosen to satisfy

$$\min_{\mathbf{r} \in \mathbb{R}^z} \mathbb{E}\left[(\Theta^{[1]} - r_L)^2\right]$$

It is easy to see that the solution to this optimisation problem is given by

$$r_l = \mathbb{E}\left[\Theta^{[1]} \mid L = l\right] = \frac{\int_0^\infty \theta^{[1]} \pi_l(\lambda\theta^{[1]}) f_{\Theta^{[1]}}(\theta^{[1]}) d\theta^{[1]}}{\int_0^\infty \pi_l(\lambda\theta^{[1]}) f_\Theta(\theta^{[1]}) d\theta^{[1]}}, \quad l = 1, \ldots, z.$$

Additionally, these relativities satisfy $\mathbb{E}[r_L] = 1$, meaning that the BMS is financially balanced once a steady state is reached.

Example 1 (The Saudi BMS). *In 2017, the Saudi Arabian Monetary Agency (SAMA) introduced a no-claim discount (NCD) reward system for the Saudi auto insurance market in response to a high number of road accidents. This system, which is a particular instance of a BMS, rewards insured drivers with discounts for consecutive years without reported claims. It operates as a* $-1/+2$ *system with six levels* ($z = 6$) *and initial class* $l^* = 6$. *The relativities associated with each level in the system are presented in Table 1.*

To quantify the MSE of the Saudi NCD system, we must first calibrate our model's parameters. For this task, we used real-life data from Allied Cooperative Insurance Group (ASIG), a Saudi auto insurance company. The dataset comprises one-year auto insurance policies recorded in 2022, with a total of 689,647 policyholders, of whom approximately 5.6% are female. The average number of claims in this portfolio is 0.0908, with a variance of 0.1553.

We then apply the method of moments to estimate the parameters describing the claim frequency λ *and* α, *obtaining* $\hat{\lambda} = 0.0908$ *and* $\hat{\alpha} = 0.1279$. *Table 1 presents the MSE for both the current relativities and the Bayesian relativities, the latter of which shows a significantly lower MSE. Note that the Bayesian relativities indicate that high penalties are required for those policyholders with high frequency. We have also included the distribution of L, which describes the distribution of policyholders within the system in the steady state. Notably, most insured people are expected to reach the level with the highest discount.*

Table 1. Current and Bayesian relativities of the Saudi BMS.

Level (l)	Current Relativities	Bayesian Relativities	$\mathbb{P}(L = l)$
6	1	10.71	0.0339
5	0.9	6.74	0.0217
4	0.8	4.95	0.0172
3	0.7	2.80	0.0324
2	0.6	2.37	0.0257
1	0.5	0.29	0.8691
MSE	7.54	3.05	

3.1.1. Number of Classes

Determining the number of levels in a BMS is a critical yet under-explored aspect of the literature. While there is extensive research on BMS design under different scenarios, there is no clear guidance on deciding the number of levels to use. Most studies either adopt arbitrary values or base their decisions on practical considerations without systematically evaluating the trade-offs between adding or subtracting levels. Here, we provide a practical and simple approach for choosing the number of levels based on the relative changes in the MSE.

Let us denote by MSE_z^{Bayes} the MSE when Bayesian relativities are used, and there are z levels in the system. The left panel of Figure 1 illustrates MSE_z^{Bayes} for z ranging from 6 to 28, with the penalty h fixed to $h = 2$ as in the Saudi BMS. We observe that as the number of classes increases, the MSE decreases, indicating improved accuracy in risk classification. However, while adding more levels enhances precision, it also introduces practical challenges. For instance, a system with too many levels may become cumbersome and less feasible to implement. Furthermore, as noted by Lemaire (1995), increasing the number of levels often results in slower convergence to the steady state, which is particularly important given that most BMSs have a limited lifespan before being revised or replaced. Therefore, fast convergence is a desirable feature of any system.

To strike a balance between accurate risk classification and a practical number of levels, we propose to analyse the absolute relative difference in MSE when additional levels are introduced. More specifically, we consider

$$\frac{\left|MSE_{z+1}^{Bayes} - MSE_{z}^{Bayes}\right|}{MSE_{z}^{Bayes}},$$

and determine the number of levels based on the above quantity being below a given threshold. The right panel of Figure 1 presents the corresponding relative differences of the MSEs, with two horizontal lines marking thresholds where changes fall below 1% (purple line) and 0.5% (orange line). Based on this analysis, we observe that if, for example, a relative difference of less than 1% is acceptable, the minimum number of classes should be set to $z = 18$. Similarly, note that if a smaller relative difference of 0.5% is required, five additional levels are needed, making $z = 23$ the minimum number of levels.

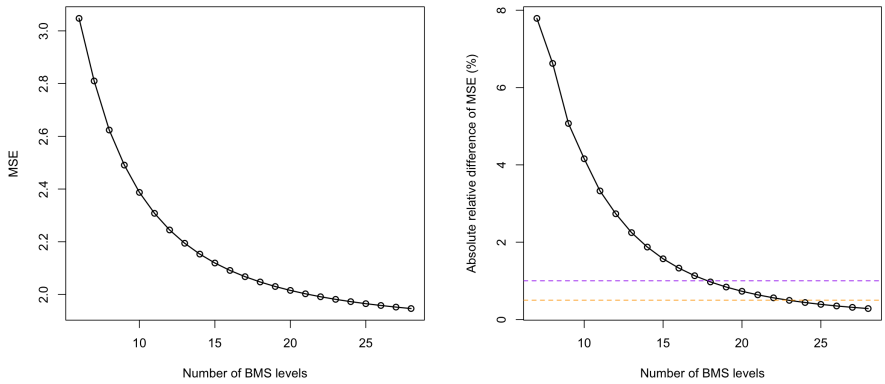


Figure 1. MSE for different numbersof levels (**left**) and their absolute relative difference (**right**).

3.1.2. Penalty of the System

In the previous analysis, we considered a fixed penalty of $h = 2$. However, it is also crucial to determine the value of h that optimises risk classification. To evaluate this component of the BMS, we calculate MSE_z^{Bayes} for different values of h and z. Specifically, we let h range from 1 to 5 and z from 6 to 28. The resulting MSE values are depicted in Figure 2, where we observe two key trends. Firstly, and as in the previous analysis, increasing the number of levels (z) leads to a more accurate risk classification. Secondly, reducing the penalty (h) lowers the MSE across all cases. In fact, $h = 1$ yields the lowest values of the MSE for all values of z.

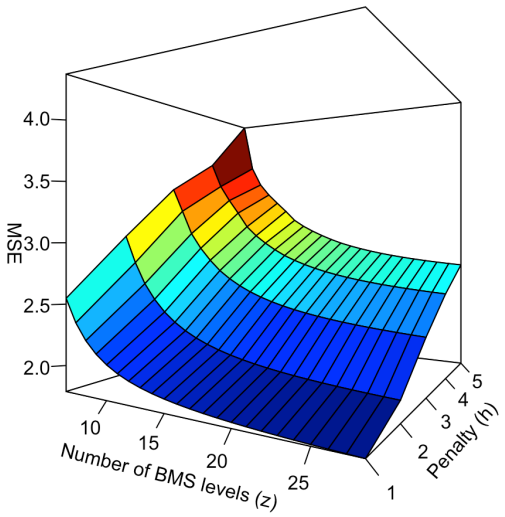


Figure 2. Surface plot of MSE for different values of z and h.

These findings are also highlighted in the top-left panel of Figure 3, where the visual representation of the trend is more evident. Note that the top-right panel of Figure 3 provides a view of the absolute relative differences in MSEs when additional levels are added for all considerations of the penalty. In particular, we observe that if a relative difference of less than 1% is required when $h = 1$, a minimum of $z = 16$ levels is necessary, while for a stricter threshold of 0.5%, a minimum of $z = 20$ levels is required. For completeness, we have included a visual representation of the MSE values for different h when $z = 16$ and $z = 20$ in the bottom panel of Figure 3. Overall, this analysis showcases the importance of selecting appropriate values for both the penalty parameter h and the number of levels z to achieve better risk classification.

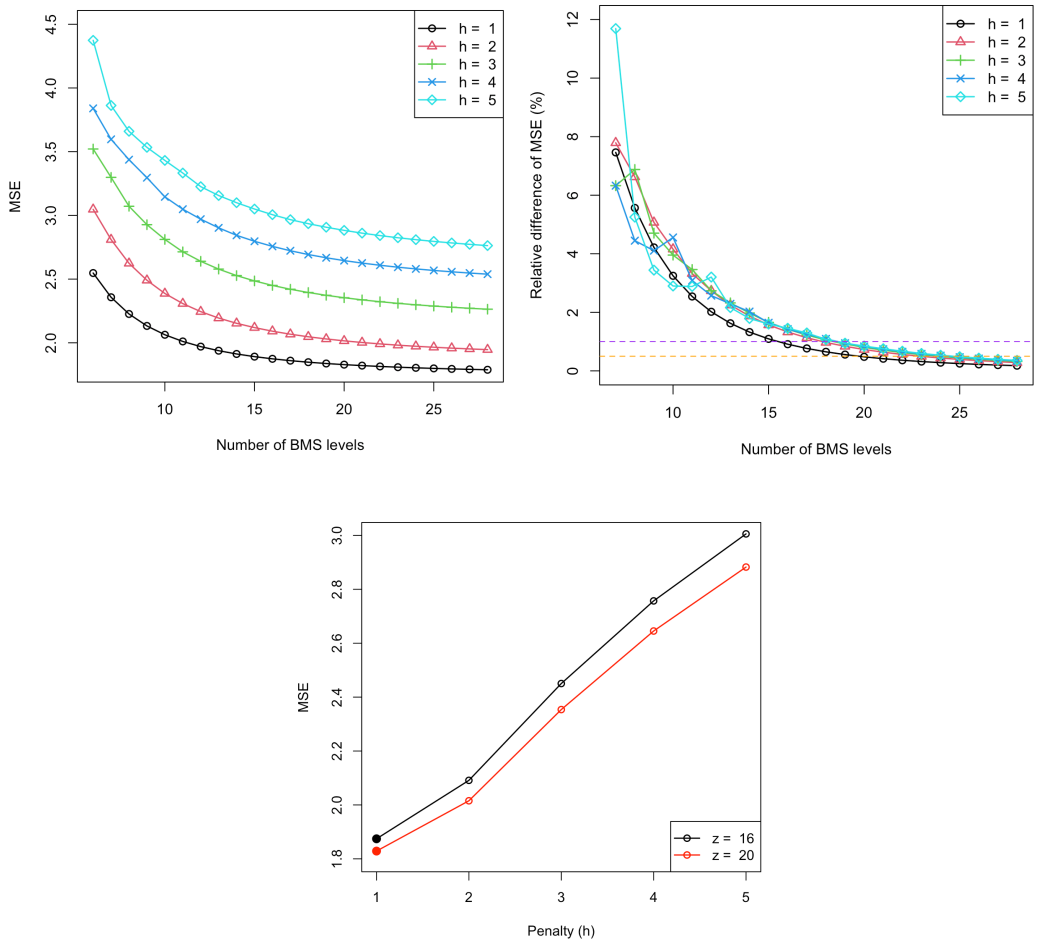


Figure 3. Contour plot of MSE for different values of z and h (**top left**) and corresponding absolute relative differences (**top right**). MSE for different values of h when $z = 16$ and $z = 20$ (**bottom**).

Remark 2 (On alternative loss functions)**.** *This study focuses on the quadratic loss function* (5)*, as it is the most commonly used measure in the literature. However, alternative loss functions can be employed, potentially leading to different results. To illustrate the effects of using an alternative loss function, we consider the exponential loss function introduced in Denuit and Dhaene* (2001) *and defined as*

$$\mathbb{E}\left[\exp\left(-c(\Theta^{[1]} - r_L)\right)\right],$$

where $c > 0$ is a "severity" parameter of the BM scale. We then compute relativities that minimise the above loss function and their corresponding loss values for different combinations of z and h when $c = 1$. Specifically, we vary z from 6 to 28 and h from 1 to 5. The results can be found in Figure 4. Specifically, the left panel shows the contour plot of the exponential loss for the different

combinations of z and h. Notably, $h = 1$ is no longer the optimal choice under this loss function across all cases. For $z \geq 12$, $h = 2$ achieves lower exponential loss values. Additionally, the right panel of Figure 4 presents the absolute relative differences of the exponential losses, where we observe that using this alternative loss function also affects the outcome of the criteria to select the number of levels. For example, if a relative difference of less than 1% is needed, a minimum of $z = 12$ levels is now required, compared to $z = 16$ under the quadratic loss function.

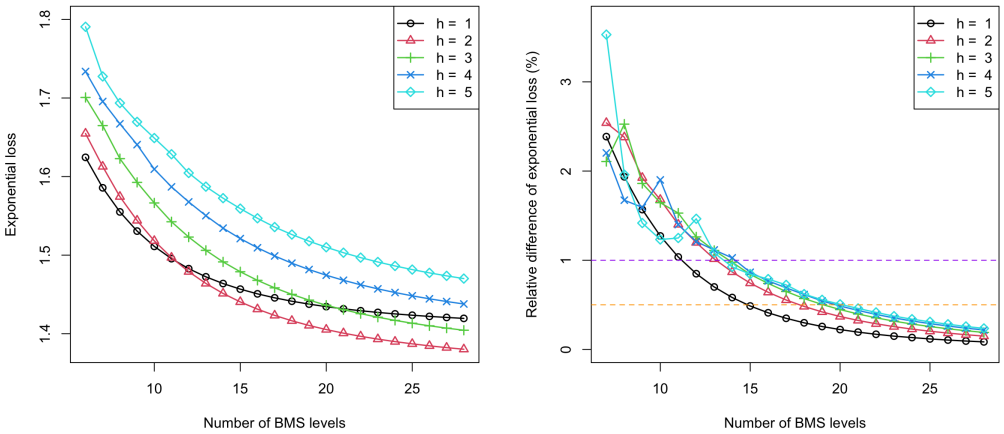


Figure 4. Contour plot of exponential loss for different values of z and h (**left**) and corresponding absolute relative differences (**right**).

3.2. Linear Relativities

One drawback of Bayesian relativities is that they may exhibit very irregular patterns, which can be undesirable for commercial purposes. To address this problem, Gilde and Sundt (1989) suggested using relativities that follow a regularly increasing pattern. More specifically, the authors suggested linear relativities of the form $r_l^{lin} := \alpha + \beta l$, $l = 1, \ldots, z$. In this way, the optimisation problem for maximum accuracy becomes

$$\min_{\alpha,\beta \in \mathbb{R}} \mathbb{E}\Big[(\Theta^{[1]} - \alpha - \beta L)^2\Big],$$

which has an explicit solution given by

$$\alpha = \mathbb{E}[\Theta^{[1]}] - \frac{\mathrm{Cov}[\Theta^{[1]}, L]}{\mathrm{Var}[L]}\mathbb{E}[L], \quad \beta = \frac{\mathrm{Cov}[\Theta^{[1]}, L]}{\mathrm{Var}[L]}.$$

The above solution yields the following expression for the linear relativities

$$r_l^{lin} = \mathbb{E}[\Theta^{[1]}] + \frac{\mathrm{Cov}[\Theta^{[1]}, L]}{\mathrm{Var}[L]}(l - \mathbb{E}[L]), \quad l = 1, \ldots, z.$$

Moreover, it is straightforward to show that, when using these linear relativities, the MSE can be computed explicitly as follows

$$\mathbb{E}\Big[(\Theta^{[1]} - \alpha - \beta L)^2\Big] = \mathrm{Var}[\Theta^{[1]}] - \frac{(\mathrm{Cov}[\Theta^{[1]}, L])^2}{\mathrm{Var}[L]}.$$

Figure 5 depicts the behaviour of the MSE for different values of z and h when linear relativities are used. As in the case of Bayesian relativities, we observe that increasing the number of levels and reducing the penalty improves accuracy in risk classification.

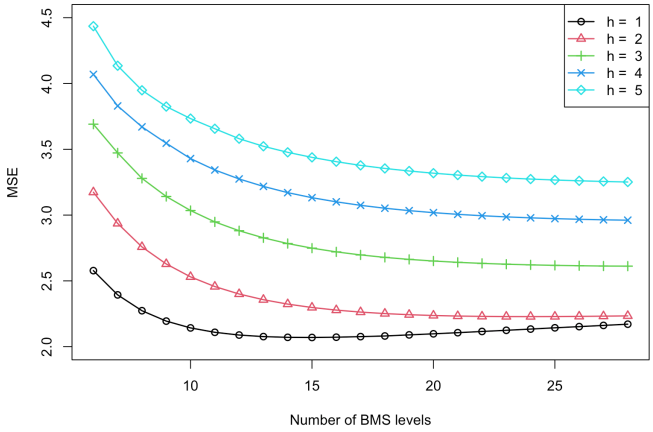


Figure 5. Contour plot of MSE for different values of z and h when linear relativities are employed.

Additionally, and similarly to the Bayesian relativities case, $h = 1$ produces the lowest MSE values across all z. Hence, we proceed to compare the behaviour of the MSE when using Bayesian and linear relativities in the case of $h = 1$. Figure 6 shows the MSE for both instances when varying z. As expected, Bayesian relativities, being the optimal solution, yield lower MSE values than linear relativities. Interestingly, the linear relativities show an optimal number of classes of $z = 15$, where a minimum value of the MSE is reached. This is in stark contrast to the Bayesian case, where we have a consistently decreasing behaviour of the MSE as the number of levels increases.

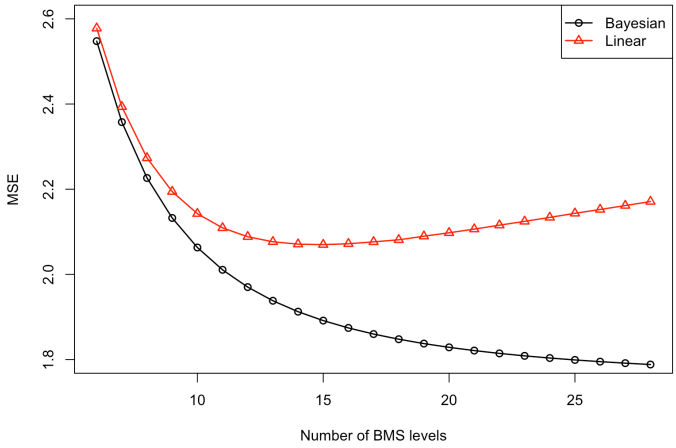


Figure 6. MSE under Bayesian and linear relativities for $-1/+1$ transition rules for different values of z.

Linear Relativities with a Fixed Initial Class

Another important component of any BMS is the initial level l^*. According to Norberg (1976), one should choose l^* such that the difference between $\mathbb{E}(S)$ and r_{l^*} is minimised, where S denotes the aggregate claim amount for a randomly selected policyholder in the steady state. Note that under our current assumptions, we have $\mathbb{E}(S) = 1$. Therefore, in an ideal situation, we should have $r_{l^*} = 1$ for some l^*. This motivates us to consider linear relativities of the form $r_l^{lin,f} := 1 + \beta(l - l^*)$, which makes $r_{l^*}^{lin,f} = 1$ for fixed l^*. With such a structure, we have that the optimisation problem translates to

$$\min_{\beta \in \mathbb{R}} \mathbb{E}\left[\left(\Theta^{[1]} - 1 - \beta(L - l^*)\right)^2\right].$$

It is easy to see that the explicit solution is given by

$$\beta = \frac{\mathrm{Cov}[\Theta^{[1]}, L]}{\mathbb{E}[(L - l^*)^2]},$$

yielding the following expression for the relativities

$$r_l^{lin,f} = 1 + \frac{\mathrm{Cov}[\Theta^{[1]}, L]}{\mathbb{E}[(L - l^*)^2]}(l - l^*), \quad l = 1, \ldots, z.$$

Moreover, the MSE can be computed explicitly via

$$\mathbb{E}\Big[(\Theta - 1 - \beta(L - l^*))^2\Big] = \mathrm{Var}[\Theta^{[1]}] - \frac{(\mathrm{Cov}[\Theta, L])^2}{\mathbb{E}[(L - l^*)^2]}.$$

From the above expression, and given that $\mathbb{E}[(L - v)^2]$, $v \in \mathbb{R}$, is minimised when $v = \mathbb{E}[L]$, it follows that the MSE with this relativity structure will always be larger compared to standard linear relativities despite the choice of l^*. Nevertheless, to minimise the MSE as much as possible, l^* should be selected as close as possible to $\mathbb{E}[L]$, which corresponds to the optimal value $l^* = \mathrm{round}\{\mathbb{E}[L]\}$, where $\mathrm{round}\{\cdot\}$ denotes rounding to an integer number. Figure 7 presents the MSE for different values of z when using the linear relativities recalibrated to have $r_{l^*} = 1$ with $l^* = \mathrm{round}\{\mathbb{E}[L]\}$ and considering $h = 1$. Not surprisingly, this structure of the relativities leads to larger values of the MSE than Bayesian and standard linear relativities. However, we still observe an optimal value for the number of levels $z = 22$, where a minimum is reached.

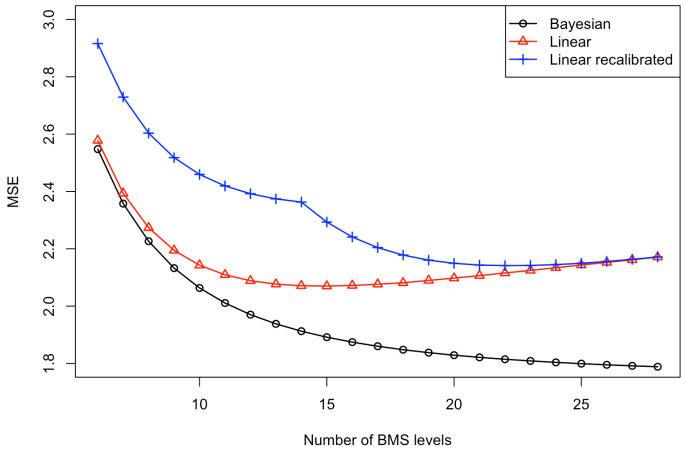


Figure 7. MSE under Bayesian, standard linear, and linear recalibrated relativities for $-1/+1$ transition rules for different values of z.

4. Transition Rules Based on Frequency and Severity

As previously noted, most implemented BMSs use only the number of reported claims to penalise a policyholder without considering the amounts of these claims. This limitation has been widely acknowledged in the literature, as large claims should intuitively be more severely penalised compared to small claims. Therefore, we now focus on studying BMS transition rules that account for both claim frequency and severity. Specifically, we consider transition rules depending on the claim sizes, separating them into two categories depending on their amount via a threshold. Although such a criterion may be somewhat unpractical due to the time needed to evaluate the cost of a claim, it still provides some valuable insights into the effects of including severity in BMS design. Furthermore, the underlying mathematical framework can easily be modified to work with another easier-

to-implement criterion, such as categorising the claims into bodily injury (BI) and property damage (PD). Since BI claims are typically larger on average than PD claims, this approach can implicitly integrate the claims' sizes into a BMS while offering a feasible pathway for implementation.

Let us denote by $\psi > 0$ a fixed and predetermined threshold such that a claim is classified as "large" if its amount exceeds ψ and "small" otherwise. Then, the number of small claims in the tth year for a random selected policyholder, denoted as N_t^I, is given by

$$N_t^I = \sum_{k=1}^{N_t} \mathbb{1}_{(Y_{t,k} \leq \psi)},$$

where $\mathbb{1}_{(\cdot)}$ denotes the indicator function. Similarly, the number of large claims, N_t^{II}, can be expressed as

$$N_t^{II} = \sum_{k=1}^{N_t} \mathbb{1}_{(Y_{t,k} > \psi)}.$$

In particular, we have that under our distributional assumption

$$\begin{aligned} N_t^I \mid (\Theta^{[1]}, \Theta^{[2]}) = (\theta^{[1]}, \theta^{[2]}) &\sim \text{Poisson}(q(\theta^{[2]})\lambda\theta^{[1]}), \\ N_t^{II} \mid (\Theta^{[1]}, \Theta^{[2]}) = (\theta^{[1]}, \theta^{[2]}) &\sim \text{Poisson}((1 - q(\theta^{[2]}))\lambda\theta^{[1]}), \end{aligned}$$

where $q(\theta^{[2]}) = \mathbb{P}(Y_{t,1} \leq \psi \mid \Theta^{[2]} = \theta^{[2]})$. Additionally, it is easy to see that the conditional joint density of the total number of claims N_t and the number of large claims N_t^{II} given $(\Theta^{[1]}, \Theta^{[2]})$ is

$$f_{N_t, N_t^{II} \mid (\Theta^{[1]}, \Theta^{[2]})}(n, m \mid (\theta^{[1]}, \theta^{[2]})) = \left(\frac{1 - q(\theta^{[2]})}{q(\theta^{[2]})} \right)^m \frac{(\lambda\theta^{[1]} q(\theta^{[2]}))^n}{m!(n-m)!} \exp(-\lambda\theta^{[1]}), \quad n, m = 0, 1, \ldots, \quad m \leq n.$$

In what follows, we focus on studying transition rules of the form $-1/+h/+h+x$, which incorporate both claim frequency and severity into the adjustment process. More specifically, under this scheme, a claim-free year results in a decrease of one level, rewarding the policyholder for having no claims. Conversely, if one or more claims are made, the level increases by $h \geq 1$ levels, representing a frequency-based penalty. Furthermore, if at least one of these claims exceeds the predefined threshold ψ, an additional penalty of $x \geq 0$ levels is applied, accounting for the severity of the claims. The precise mathematical definition of this scheme is given by

$$L_{t+1} = \begin{cases} \max\{L_t - 1, 1\}, & N_t = 0, \\ \min\{L_t + h, z\}, & N_t > 0, N_t^{II} = 0, \\ \min\{L_t + h + x, z\}, & N_t > 0, N_t^{II} > 0. \end{cases}$$

We illustrate how the $-1/+h/+h+x$ transition rules work by considering a system with $z = 10$ levels, $h = 2$, and $x = 2$, meaning a $-1/+2/+4$ system. Consider a policyholder that is presently at level $L_t = 5$. If the policyholder has no claims during the current policy year ($N_t = 0$), then the level for the next year is $L_{t+1} = 4$. Conversely, if the policyholder has at least one claim and all of them are below the predefined threshold ψ ($N_t > 0$, $N_t^{II} = 0$), then their level for the following policy year is $L_{t+1} = 7$. Finally, if the policyholder presents at least one claim above the threshold ψ ($N_t > 0$, $N_t^{II} > 0$), then their level for the next year is $L_{t+1} = 9$.

Note that under this structure of BMS's rules, the non-zero elements of the one-step transition matrix for a policyholder with an expected claim frequency $\lambda\theta^{[1]}$ and severity $\theta^{[2]}$ can be computed using the following expressions:

$$
\begin{aligned}
&p_{i,\max(i-1,1)}(\lambda\theta^{[1]},\theta^{[2]}) = s_0(\lambda\theta^{[1]},\theta^{[2]}), \quad i = 1,\dots,z,\\
&p_{i,\min(i+h,z)}(\lambda\theta^{[1]},\theta^{[2]}) = 1 - s_0(\lambda\theta^{[1]},\theta^{[2]}), \quad i = z-h,\dots,z,\\
&p_{i,\min(i+h,z)}(\lambda\theta^{[1]},\theta^{[2]}) = s_{\geq 1,0}(\lambda\theta^{[1]},\theta^{[2]}), \quad i = 1,\dots,z-h-1,\\
&p_{i,\min(i+h+x,z)}(\lambda\theta^{[1]},\theta^{[2]}) = 1 - s_0(\lambda\theta^{[1]},\theta^{[2]}) - s_{\geq 1,0}(\lambda\theta^{[1]},\theta^{[2]}), \quad i = 1,\dots,z-h-1,
\end{aligned}
$$

where

$$
\begin{aligned}
s_0(\lambda\theta^{[1]},\theta^{[2]}) &= \mathbb{P}(N_t = 0 \mid (\Theta^{[1]},\Theta^{[2]}) = (\theta^{[1]},\theta^{[2]}))\\
&= \exp(-\lambda\theta^{[1]})
\end{aligned}
$$

and

$$
\begin{aligned}
s_{\geq 1,0}(\lambda\theta^{[1]},\theta^{[2]}) &= \mathbb{P}(N_t \geq 1, N_t^{II} = 0 \mid (\Theta^{[1]},\Theta^{[2]}) = (\theta^{[1]},\theta^{[2]}))\\
&= \exp(-\lambda\theta^{[1]}(1 - q(\theta^{[2]}))) - \exp(-\lambda\theta^{[1]}).
\end{aligned}
$$

Note that the above expressions are valid when $x \geq 1$ and that the case of $x = 0$ can be computed using the corresponding formulas in Section 3. In fact, the special instance of $x = 0$ corresponds to the previously studied case of frequency-dependent transition rules, as no distinction between claim sizes is made.

With a structure of the transition rules established, we now proceed to the mathematical foundations involved in the calculation of Bayesian and linear relativities under the present setting. Starting with the Bayesian relativities, these should be chosen such that they minimise the MSE, that is,

$$
\min_{\mathbf{r}\in\mathbb{R}^z} \mathbb{E}\left[\left(\Theta^{[1]}\Theta^{[2]} - r_L\right)^2\right].
$$

It is easy to see that the solution to this optimisation problem is given by

$$
r_l = \mathbb{E}\left[\Theta^{[1]}\Theta^{[2]} \mid L = l\right] = \frac{\int_0^\infty \int_0^\infty \theta^{[1]}\theta^{[2]}\pi_l(\lambda\theta^{[1]},\theta^{[2]}) f_{\Theta^{[1]},\Theta^{[2]}}(\theta^{[1]},\theta^{[2]}) d\theta^{[1]} d\theta^{[2]}}{\int_0^\infty \int_0^\infty \pi_l(\lambda\theta^{[1]},\theta^{[2]}) f_{\Theta^{[1]},\Theta^{[2]}}(\theta^{[1]},\theta^{[2]}) d\theta^{[1]} d\theta^{[2]}}, \quad l = 1,\dots,z.
$$

For the case of linear relativities of the form $r_l^{lin} = \alpha + \beta l$, $l = 1,\dots,z$, we have that the optimisation problem translates to

$$
\min_{\alpha,\beta\in\mathbb{R}} \mathbb{E}\left[\left(\Theta^{[1]}\Theta^{[2]} - \alpha - \beta L\right)^2\right].
$$

In an analogous way to the case of frequency-dependent rules (see Gilde and Sundt (1989) for details), one can show that the explicit solution for α and β are given by

$$
\alpha = \mathbb{E}[\Theta^{[1]}\Theta^{[2]}] - \frac{\mathrm{Cov}[\Theta^{[1]}\Theta^{[2]}, L]}{\mathrm{Var}[L]}\mathbb{E}[L], \quad \beta = \frac{\mathrm{Cov}[\Theta^{[1]}\Theta^{[2]}, L]}{\mathrm{Var}[L]},
$$

implying

$$
r_l^{lin} = \mathbb{E}[\Theta^{[1]}\Theta^{[2]}] + \frac{\mathrm{Cov}[\Theta^{[1]}\Theta^{[2]}, L]}{\mathrm{Var}[L]}(l - \mathbb{E}[L]), \quad l = 1,\dots,z.
$$

Moreover, when using these relativities, the MSE can be computed via

$$\mathbb{E}\left[\left(\Theta^{[1]}\Theta^{[2]} - \alpha - \beta L\right)^2\right] = \mathrm{Var}[\Theta^{[1]}\Theta^{[2]}] - \frac{(\mathrm{Cov}[\Theta^{[1]}\Theta^{[2]}, L])^2}{\mathrm{Var}[L]}$$

We finally note that when considering linear relativities with a fixed initial class l^* such that $r_{l^*} = 1$, that is, of the form $r_l^{lin,f} = 1 + \beta(l - l^*)$, we have the optimisation problem

$$\min_{\beta \in \mathbb{R}} \mathbb{E}\left[\left(\Theta^{[1]}\Theta^{[2]} - 1 - \beta(L - l^*)\right)^2\right].$$

As in the case of standard linear relativities, there exists an explicit solution for β given by

$$\beta = \frac{\mathrm{Cov}[\Theta^{[1]}\Theta^{[2]}, L] + (\mathbb{E}[L] - l^*)(\mathbb{E}[\Theta^{[1]}\Theta^{[2]}] - 1)}{\mathbb{E}[(L - l^*)^2]},$$

which yields

$$r_l^{lin,f} = 1 + \frac{\mathrm{Cov}[\Theta^{[1]}\Theta^{[2]}, L] + (\mathbb{E}[L] - l^*)(\mathbb{E}[\Theta^{[1]}\Theta^{[2]}] - 1)}{\mathbb{E}[(L - l^*)^2]}(l - l^*), \quad l = 1, \ldots, z.$$

Additionally, the MSE can be computed through the following simplified formula:

$$\mathbb{E}\left[\left(\Theta^{[1]}\Theta^{[2]} - 1 - \beta(L - l^*)\right)^2\right] = \mathrm{Var}[\Theta^{[1]}\Theta^{[2]}] + \left(\mathbb{E}[\Theta^{[1]}\Theta^{[2]}] - 1\right)^2 - \frac{(\mathrm{Cov}[\Theta^{[1]}\Theta^{[2]}, L] + (\mathbb{E}[L] - l^*)(\mathbb{E}[\Theta^{[1]}\Theta^{[2]}] - 1))^2}{\mathbb{E}[(L - l^*)^2]}.$$

From this last expression, it is evident that, when working with this form of relativities, l^* should be chosen as close as possible to $\mathbb{E}[L]$ to obtain the minimum value of the MSE.

4.1. Application to the Saudi Arabian Auto Insurance Market

We now aim to explore the impact of the distinct components of a BMS on the MSE under different intensity structures within the present scheme of frequency–severity transition rules. To this end, we begin by calibrating the parameter σ, which characterises the severity distribution, using the method of moments. After scaling the ASIG dataset to have a sample mean claim size of 1, we estimate $\hat{\sigma} = 0.9057$. Inspired by the proportion of BI claims in 2019 in Saudi Arabia of approximately 17% among all auto insurance claims (c.f. Ministry of Health 2021), we set $\psi = 1.8828$, so that $\mathbb{P}(Y_{t,1} \leq \psi) = 0.83$. Additionally, to delve into the effect of the dependence structure of $(\Theta^{[1]}, \Theta^{[2]})$ into the results, we consider two scenarios: one assumes independence, while the other introduces a negative correlation. The reasoning behind the first case is that this is a prevalent assumption within the literature. The second case builds on existing research indicating that women, while potentially involved in a higher frequency of accidents, tend to be associated with less severe incidents compared to men (see, e.g., Al-Garawi et al. 2021). This choice of negative correlation is further supported by recent studies such as Garrido et al. (2016); Oh et al. (2020), which, using real-world data, provide evidence of a negative association between frequency and severity in auto insurance.

4.1.1. Independent Random Effects

We start by considering the case of $\Theta^{[1]}$ independent of $\Theta^{[2]}$. As in the previous section, our focus is on studying the impact of three key features of the BMS in the MSE: the number of levels, the transition rules, and the structure of the relativities. We begin the analysis

by considering Bayesian relativities with a fixed $z = 16$ number of levels. The choice of $z = 16$ is due to the fact that this leads to an absolute relative difference of less than 1% in the case of frequency-dependent rules when $h = 1$. Moreover, note that the assumption of $\Theta^{[1]} \perp\!\!\!\perp \Theta^{[2]}$ implies the following simplified form of the Bayesian relativities

$$r_l = \frac{\int_0^\infty \int_0^\infty \theta^{[1]}\theta^{[2]}\pi_l(\lambda\theta^{[1]}, \theta^{[2]}) f_{\Theta^{[1]}}(\theta^{[1]}) f_{\Theta^{[2]}}(\theta^{[2]}) d\theta^{[1]} d\theta^{[2]}}{\int_0^\infty \int_0^\infty \pi_l(\lambda\theta^{[1]}, \theta^{[2]}) f_{\Theta^{[1]}}(\theta^{[1]}) f_{\Theta^{[2]}}(\theta^{[2]}) d\theta^{[1]} d\theta^{[2]}}, \quad l = 1, \ldots, z.$$

We are now interested in determining the values of h and x that optimise risk classification. To this end, we calculate the MSE for h ranging from 1 to 5 and x ranging from 0 to 5. The resulting behaviour of the MSE with respect to these transition rule parameters is depicted in Figure 8. In particular, we observe that $h = 1$ consistently yields the lowest MSE values. Additionally, we note that adding a penalty for large claim sizes, that is, $x > 0$, improves risk classification. Moreover, we identify optimal values of $h = 1$ and $x = 3$, where the MSE reaches a minimum value. Overall, we see that the selection of both components h and x plays a crucial role in the effectiveness of the BMS for risk classification.

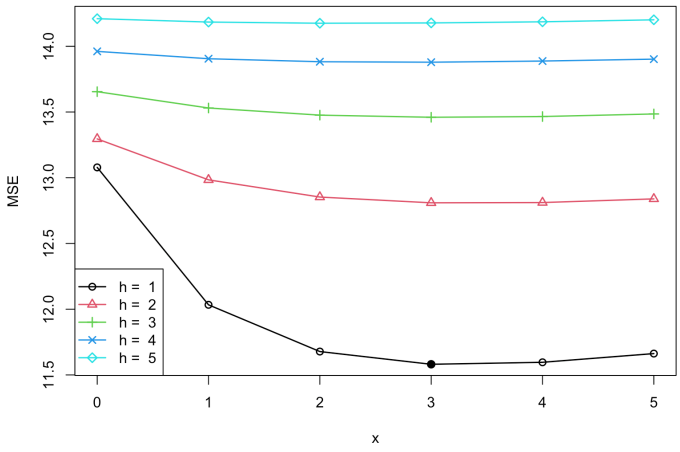


Figure 8. MSE for different values of h and x with $z = 16$.

We proceed by computing the MSE for various combinations of z, h, and x to explore the effect of the number of levels on risk classification. Specifically, we consider z from 6 to 28, and once again, h from 1 to 5 and x from 0 to 5. For each value of z, we identify the combination of h and x that leads to the minimum value of the MSE. The optimal values of h and x are depicted in the top panel of Figure 9. In particular, we observe that $h = 1$ is optimal across all cases, while x increases with respect to the number of levels (z). The corresponding MSE values for these optimal combinations are shown in the bottom left panel of Figure 9, where we observe that increasing the number of levels consistently improves risk classification. Additionally, we have included the absolute relative changes on the MSEs in the bottom right panel of Figure 9. These changes can be used to determine the minimum number of levels required to meet a specified tolerance threshold. For instance, for a tolerance level of 1%, $z = 12$ is sufficient (with $h = 1$, $x = 3$), while for a stricter threshold of 0.5%, $z = 15$ (and $h = 1$, $x = 3$) is necessary.

We now consider linear relativities in the analysis. Firstly, note that the assumption of independence on the random effects yields the following simplified formula for the standard linear relativities:

$$r_l^{lin} = 1 + \frac{\mathrm{Cov}[\Theta^{[1]}\Theta^{[2]}, L]}{\mathrm{Var}[L]}(l - \mathbb{E}[L]), \quad l = 1, \ldots, z.$$

Moreover, the corresponding MSE can be computed as

$$\mathbb{E}\left[\left(\Theta^{[1]}\Theta^{[2]} - \alpha - \beta L\right)^2\right] = \left(\mathbb{E}[\Theta^{[1]}]\right)^2 \mathrm{Var}[\Theta^{[2]}] + \mathrm{Var}[\Theta^{[1]}] - \frac{(\mathrm{Cov}[\Theta^{[1]}\Theta^{[2]}, L])^2}{\mathrm{Var}[L]}$$

We now calculate the MSE when linear relativities are used, considering the same combinations of z, h, and x values as in the Bayesian case. Interestingly, the combination of values of h and x leading to the lowest values of the MSE for each z differs from the results obtained with Bayesian relativities (see the left panel of Figure 10). Specifically, although we still observe a $h = 1$ optimal in all cases of z, we have a distinct behaviour of x. The right panel of Figure 10 compares the MSE obtained using Bayesian and linear relativities, each calculated with their respective optimal values of h and x. This comparison highlights the differences in performance between the two structures of relativities.

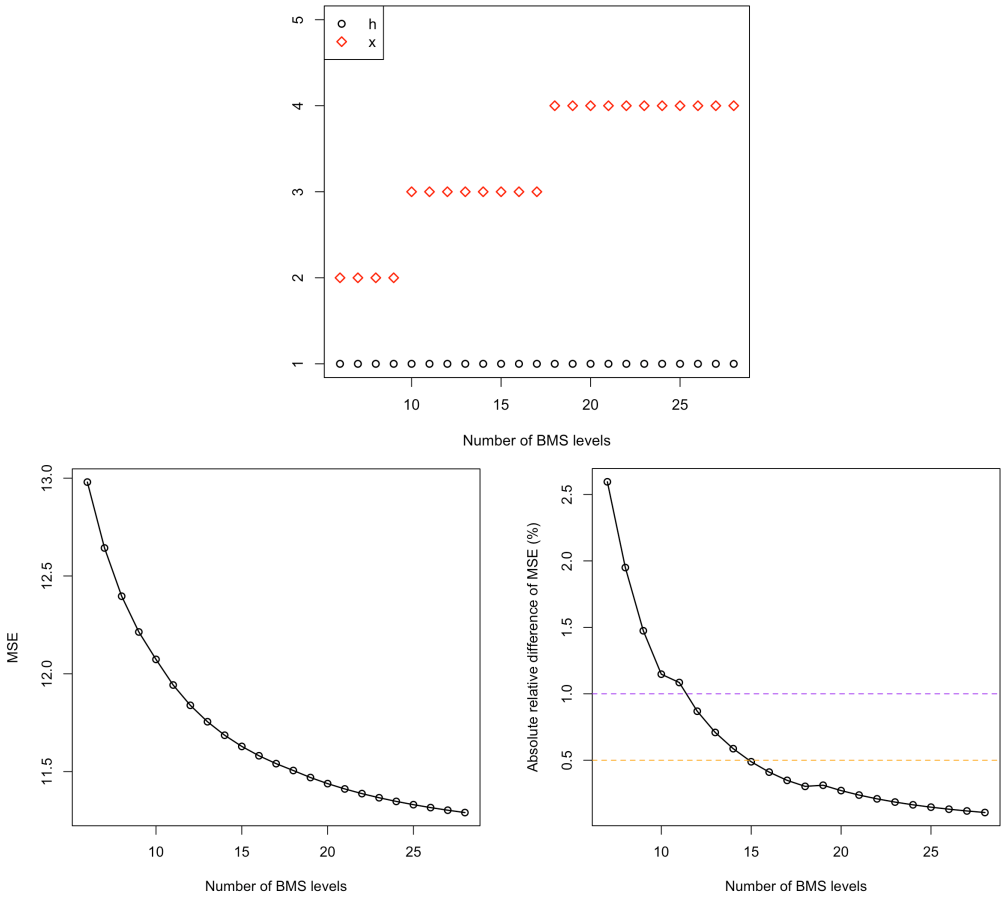


Figure 9. Values of h and x that minimise the MSE for each value of z (**top**), corresponding MSE values (**bottom left**), and their absolute relative differences (**bottom right**).

To conclude the analysis of the independent case, we proceed to recalibrate the linear relativities to have an initial level l^* such that $r_{l^*} = 1$. We first note that their computation simplifies to the following formula

$$r_l^{lin,f} = 1 + \frac{\mathrm{Cov}[\Theta^{[1]}\Theta^{[2]}, L]}{\mathbb{E}[(L - l^*)^2]}(l - l^*), \quad l = 1, \ldots, z.$$

with the corresponding MSE given explicitly by

$$\mathbb{E}\left[\left(\Theta^{[1]}\Theta^{[2]} - 1 - \beta(L - l^*)\right)^2\right] = \left(\mathbb{E}[\Theta^{[1]}]\right)^2 \mathrm{Var}[\Theta^{[2]}] + \mathrm{Var}[\Theta^{[1]}] - \frac{(\mathrm{Cov}[\Theta^{[1]}\Theta^{[2]}, L])^2}{\mathbb{E}[(L - l^*)^2]}.$$

Figure 11 illustrates the resulting MSE values, where, for each z, we have employed the optimal values of h and x identified for the standard linear relativities (shown in Figure 10), along with $l^* = \text{round}\{\mathbb{E}[L]\}$. This highlights the potential cost of allowing this restriction in the structure of linear relativities.

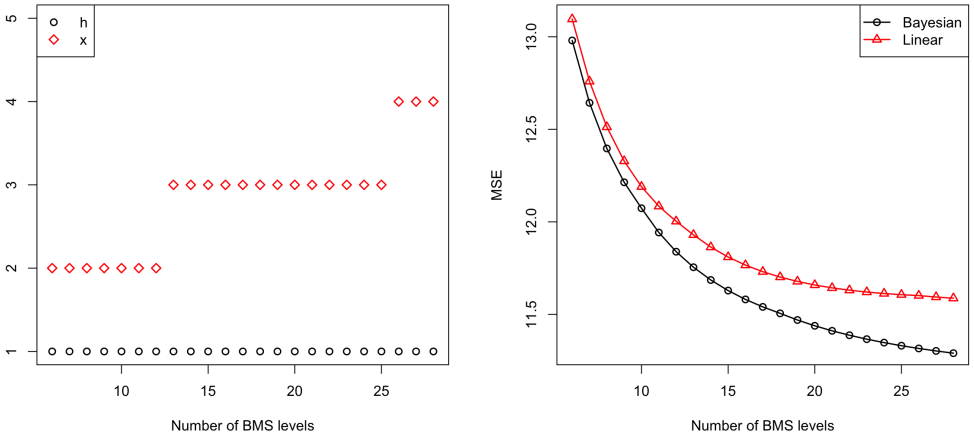


Figure 10. Values of h and x that minimise the MSE for each value of z when standard linear relativities are employed (**left**), and corresponding MSE values (**right**).

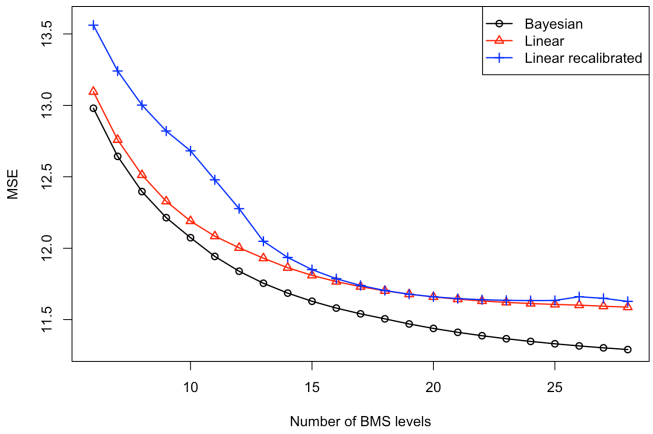


Figure 11. MSE under Bayesian, standard linear, and linear recalibrated relativities for different values of z.

4.1.2. Negatively Correlated Random Effects

We now consider the case of negatively correlated random effects. More specifically, and for illustration purposes, we assume that the joint distribution function $F_{\Theta^{[1]},\Theta^{[2]}}$ of $(\Theta^{[1]}, \Theta^{[2]})$ is given by

$$F_{\Theta^{[1]},\Theta^{[2]}}(\theta^{[1]}, \theta^{[2]}) = C_\rho(F_{\Theta^{[1]}}(\theta^{[1]}), F_{\Theta^{[2]}}(\theta^{[2]}))$$

where C_ρ denotes a bivariate Gaussian copula with correlation parameter ρ. For this analysis, we set $\rho = -0.5$, so that the corresponding Kendall's tau ρ_τ is $\rho_\tau = -1/3$. While this parameter choice is somewhat arbitrary and, in practice, it should be calibrated to reflect the data at hand, it will serve as a proof of concept to demonstrate the impact of negative correlation in risk classification.

As in the independent case, we begin by setting $z = 16$ and compute the MSE for different combinations of values of h and x. Specifically, h is varied from 1 to 5 and x from 0 to 10. The resulting values of the MSE are displayed in Figure 12. We observe that incorporating a penalty for large claims, that is, $x > 0$, consistently reduces the MSE.

Interestingly, while $h = 1$ yields significantly lower values of the MSE for values of $x > 2$, this is not the case for $x < 2$. This finding contrasts with the previous instance, where $h = 1$ always produced the lowest MSE values regardless of x. An overall minimum is observed at the combination $h = 1$, $x = 7$, indicating a substantially bigger penalty for large claims compared to the independent case. This suggests that the dependence structure has a substantial influence on the results.

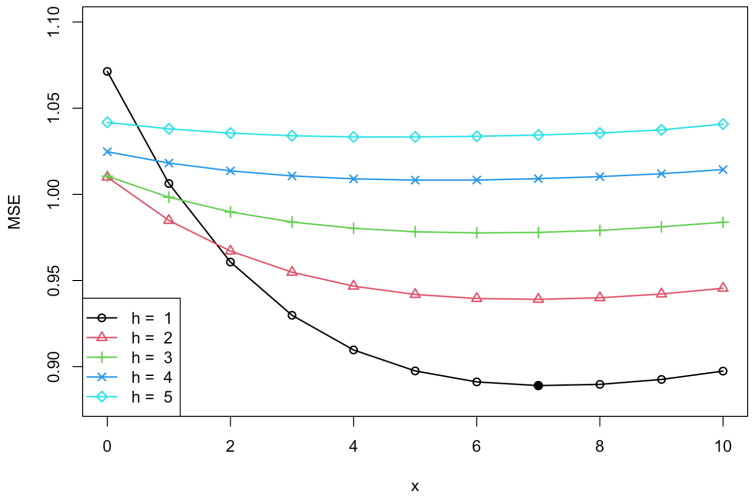


Figure 12. MSE for different values of h and x with $z = 16$ for the negatively correlated case.

We continue the analysis by computing the MSE for different combinations of z, h, and x, where we range z from 6 to 28, h from 1 to 5, and x from 0 to 12. The results are summarised in Figure 13. The top left panel shows the optimal combination of h and x that minimises the MSE for each z, while the top right panel displays the corresponding MSE values. Note that, as in the independent case, $h = 1$ is optimal across all cases; however, more substantial penalties x that increase with z are now required. For completeness, we have also included the absolute relative differences of these MSE values in the bottom left panel of Figure 9.

To provide insight into the potential effects of computing relativities under a misspecified dependence assumption, we evaluate the performance of Bayesian relativities derived under the assumption of independence when applied to the current scenario of negatively correlated random effects. Specifically, we use the Bayesian relativities computed as if $\Theta^{[1]}$ and $\Theta^{[2]}$ were independent (for the combinations of h and x values in Figure 12) and recalculate the corresponding MSE, accounting for the negative correlation between $\Theta^{[1]}$ and $\Theta^{[2]}$. The results can be found in Figure 14, where we also include the MSE values computed using Bayesian relativities that properly account for the dependence structure. A significant difference is observed between the two curves, reflecting the potential adverse impact of using relativities derived under an incorrect independence assumption.

We complement the previous analysis by considering the specific case of $x = 0$, corresponding to a system with no differentiation based on claim sizes. Specifically, we take the relativities obtained in Section 3 for the $-1/+1$ transition rules (which assume independence of the random effects) and compute the MSE under the negative correlation assumption. This scenario allows us to explore further the implications of using Bayesian relativities derived under a misspecified dependence assumption. Similar to the findings in Oh et al. (2020), and as depicted in Figure 15, we observe that failing to account for the dependence between $\Theta^{[1]}$ and $\Theta^{[2]}$ leads to consistently larger MSE values. Overall, these two previously scenarios highlight the importance of accurately modelling the dependence between random effects to achieve optimal risk classification.

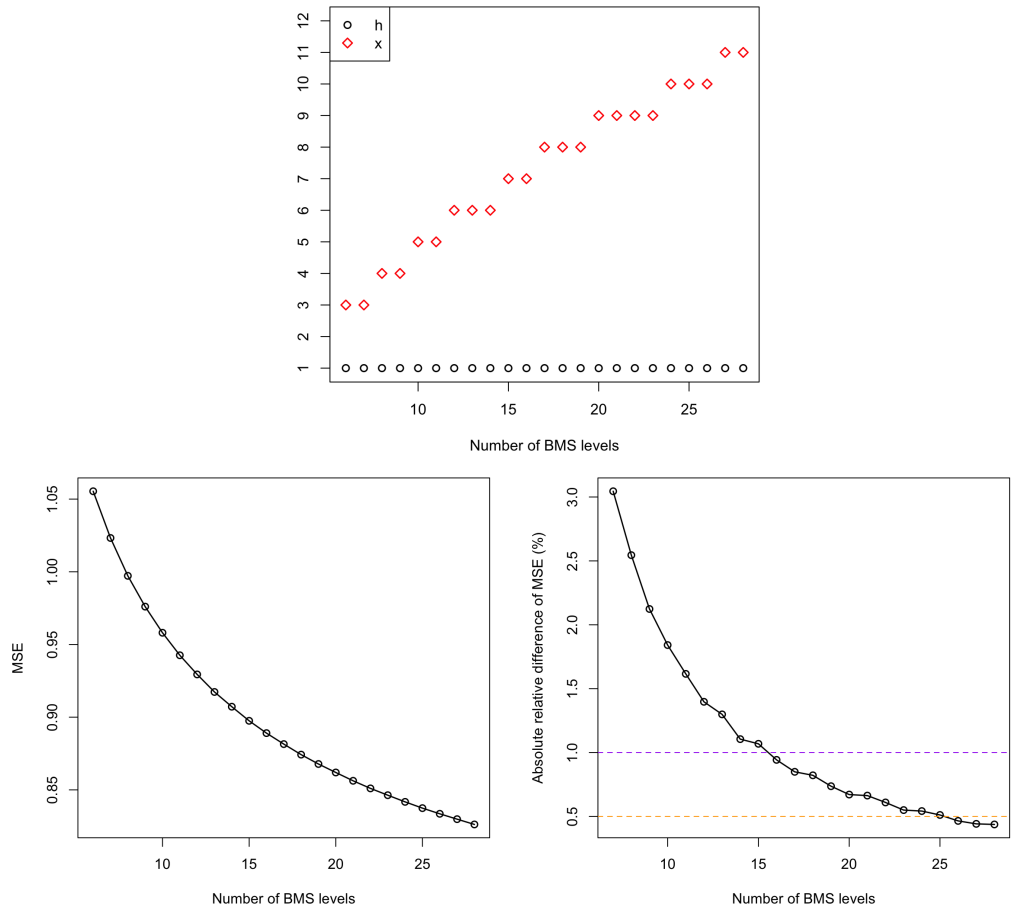


Figure 13. Values of h and x that minimise the MSE for each value of z under negative correlation (**top**), corresponding MSE values (**bottom left**), and their absolute relative differences (**bottom right**).

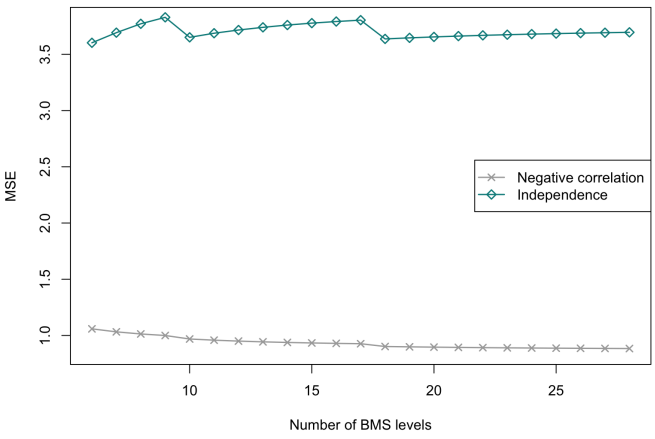


Figure 14. MSE values with relativities computed under the assumption of independence and those accounting for negative correlation, evaluated at the optimal values of h and x identified in the independent case.

We conclude the study by determining the optimal values of h and x for all z between 6 and 28 when using standard linear relativities. The results, including the corresponding MSE values, are presented in Figure 16. Similar to the independent case, we observe a slight difference in the combination of values of h and x compared to the Bayesian case. Additionally, we have included the MSE curve obtained using Bayesian relativities to illustrate the performance cost of changing the relativities' structure.

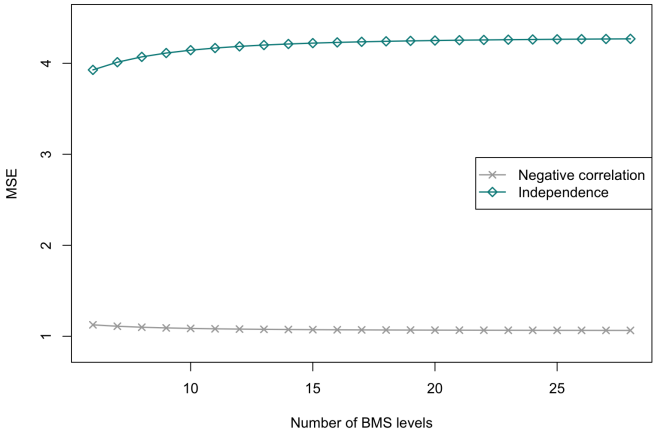


Figure 15. MSE values with relativities computed under the assumption of independence and those accounting for negative correlation when penalties only consider the claim frequency.

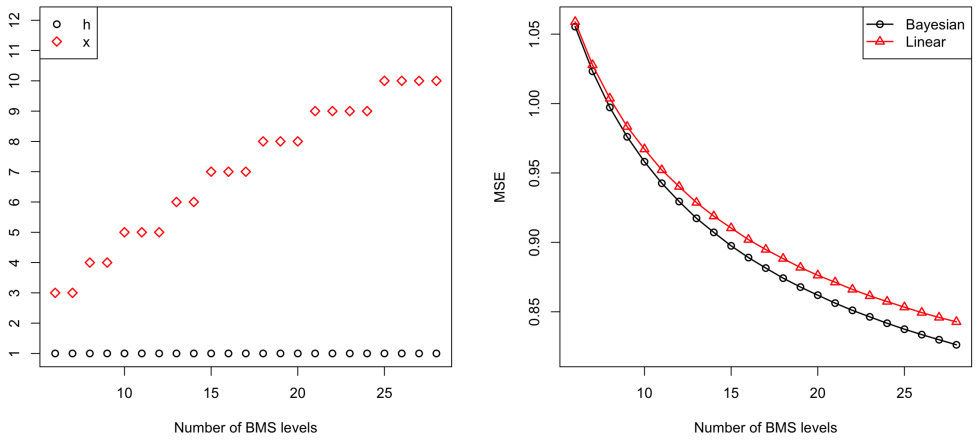


Figure 16. Values of h and x that minimise the MSE for each value of z when standard linear relativities are employed (**left**), and corresponding MSE values (**right**).

5. Conclusions

Most studies regarding Bonus–Malus system (BMS) design assume fixed system characteristics, such as predefined transition rules and a set number of levels, without fully exploring the impact of these choices on the results. In this paper, we have presented a comprehensive analysis of how modifications to key BMS components affect effective risk classification, measured in terms of the mean square error (MSE). More specifically, we show that increasing the number of levels generally improves classification accuracy. Nevertheless, to maintain a practical balance between the number of levels and the reduction in MSE, we proposed evaluating the relative changes in MSE with respect to the number of levels. Then, one can select an optimal number of levels for which these changes fall below a specified threshold, thus ensuring efficiency and practicality. Additionally, our results highlight that including penalties that depend not only on the number of claims but also on the claim sizes significantly enhances fairness and that careful consideration of the selection of these penalties is crucial for maximum accuracy.

This paper also explored the influence of the dependence between the claim frequency and severity on the results. We found that the dependence structure significantly affects the optimal parameter values and the resulting MSE. Notably, the use of relativities derived under a misspecified dependence assumption leads to suboptimal performance, emphasising the need for accurate modelling of dependence structures.

Overall, this study offers actionable insights for designing fair insurance pricing schemes, which can be applied to markets with a pressing need for fairer systems, such as the post-decree Saudi Arabian car insurance market.

Author Contributions: Conceptualization, A.A., C.C and J.Y.; methodology, A.A., C.C and J.Y.; software, A.A. and J.Y.; validation, A.A., C.C and J.Y.; formal analysis, A.A., C.C and J.Y.; investigation, A.A., C.C and J.Y.; resources, A.A., C.C and J.Y.; data curation, A.A.; writing—original draft preparation, A.A. and J.Y; writing—review and editing, A.A., C.C and J.Y.; visualization, A.A. and J.Y.; supervision, C.C. and J.Y.; project administration, C.C.; funding acquisition, A.A. All authors have read and agreed to the published version of the manuscript.

Funding: Asrar Alyafie's research is financially supported by the Ministry of Education of the Kingdom of Saudi Arabia, represented by the Saudi Arabian Cultural Bureau (SACB), and by the University of Jeddah.

Data Availability Statement: Restrictions apply to the availability of these data. Data were obtained from Allied Cooperative Insurance Group and are available from the authors with the permission of Allied Cooperative Insurance Group.

Acknowledgments: We want to thank Allied Cooperative Insurance Group for providing us access to the mentioned dataset and especially Mohammed A. Al-Gadhi for sharing his expertise in the Saudi Arabian insurance market.

Conflicts of Interest: The authors declare no conflicts of interest.

References

Afonso, Lourdes B., Rui M. R. Cardoso, Alfredo D. Egídio dos Reis, and Gracinda Rita Guerreiro. 2017. Measuring the impact of a Bonus–Malus system in finite and continuous time ruin probabilities for large portfolios in motor insurance. *ASTIN Bulletin: The Journal of the IAA* 47: 417–35. [CrossRef]

Ahn, Jae Youn, Eric C. K. Cheung, Rosy Oh, and Jae-Kyung Woo. 2022. Optimal relativities in a modified Bonus–Malus system with long memory transition rules and frequency–severity dependence. *Variance* 15.

Al-Garawi, Najah, Muhammad Abubakar Dalhat, and Omer Aga. 2021. Assessing the road traffic crashes among novice female drivers in Saudi Arabia. *Sustainability* 13: 8613. [CrossRef]

Alyafie, Asrar, Corina Constantinescu, and Jorge Yslas. 2023. An analysis of the current Saudi Arabian no-claim discount system and its adaptability for novice women drivers. *CAS E-Forum* Spring.

Asmussen, Søren, Corina Constantinescu, and Julie Thøgersen. 2021. On the risk of credibility premium rules. *Scandinavian Actuarial Journal* 2021: 866–89. [CrossRef]

Denuit, Michel, and Jan Dhaene. 2001. Bonus–Malus scales using exponential loss functions. *Blätter der DGVFM* 25: 13–27. [CrossRef]

Denuit, Michel, Xavier Maréchal, Sandra Pitrebois, and Jean-François Walhin. 2007. *Actuarial Modelling of Claim Counts: Risk Classification, Credibility and Bonus–Malus Systems*. Hoboken: John Wiley & Sons.

Dionne, Georges, and Olfa Ghali. 2005. The (1992) Bonus–Malus system in tunisia: An empirical evaluation. *Journal of Risk and Insurance* 72: 609–33. [CrossRef]

Frangos, Nicholas E., and Spyridon D. Vrontos. 2001. Design of optimal Bonus–Malus systems with a frequency and a severity component on an individual basis in automobile insurance. *ASTIN Bulletin: The Journal of the IAA* 31: 1–22. [CrossRef]

Frees, Edward W., Gee Lee, and Lu Yang. 2016. Multivariate frequency–severity regression models in insurance. *Risks* 4: 4. [CrossRef]

Garrido, José, Christian Genest, and Juliana Schulz. 2016. Generalized linear models for dependent frequency and severity of insurance claims. *Insurance: Mathematics and Economics* 70: 205–15. [CrossRef]

Gilde, Vibeke, and Bjørn Sundt. 1989. On bonus systems with credibility scales. *Scandinavian Actuarial Journal* 1989: 13–22. [CrossRef]

Gómez-Déniz, Emilio. 2016. Bivariate credibility Bonus–Malus premiums distinguishing between two types of claims. *Insurance: Mathematics and Economics* 70: 117–24. [CrossRef]

Heras Angeles, José Luis Vilar Zanón, and José Antonio Gil Fana. 2002. Asymptotic fairness of Bonus–Malus systems and optimal scales of premiums. *The Geneva Papers on Risk and Insurance Theory* 27: 61–82. [CrossRef]

Heras, Antonio, José A. Gil, Pilar García-Pineda, and José L. Vilar. 2004. An application of linear programming to bonus malus system design. *ASTIN Bulletin: The Journal of the IAA* 34: 435–56. [CrossRef]

Lemaire, Jean. 1995. *Bonus–Malus Systems in Automobile Insurance*. Berlin/Heidelberg: Springer Science & Business Media, vol. 19.

Li, Bo, Weihong Ni, and Corina Constantinescu. 2015. Risk models with premiums adjusted to claims number. *Insurance: Mathematics and Economics* 65: 94–102. [CrossRef]

Ministry of Health. 2021. MOH Statistics and Indicators: Road Traffic Injuries and Deaths. Available online: https://www.moh.gov.sa/en/Ministry/Statistics/Pages/Traffic-accidents.aspx (accessed on 8 May 2022).

Moreno, Ignacio, Francisco J. Vázquez, and Richard Watt. 2006. Can Bonus–Malus allieviate insurance fraud? *Journal of Risk and Insurance* 73: 123–51. [CrossRef]

Ni, Weihong, Corina Constantinescu, and Athanasios A. Pantelous. 2014. Bonus–malus systems with Weibull distributed claim severities. *Annals of Actuarial Science* 8: 217–33. [CrossRef]

Norberg, Ragnar. 1976. A credibility theory for automobile bonus systems. *Scandinavian Actuarial Journal* 1976: 92–107. [CrossRef]

Oh, Rosy, Joseph H. T. Kim, and Jae Youn Ahn. 2022. Designing a Bonus–Malus system reflecting the claim size under the dependent frequency–severity model. *Probability in the Engineering and Informational Sciences* 36: 963–87. [CrossRef]

Oh, Rosy, Peng Shi, and Jae Youn Ahn. 2020. Bonus–Malus premiums under the dependent frequency–severity modeling. *Scandinavian Actuarial Journal* 2020: 172. [CrossRef]

Pitrebois, Sandra, Michel Denuit, and Jean-François Walhin. 2003a. Marketing and Bonus–Malus systems. Paper presented at the 2003 ASTIN Colloquium, Berlin, Germany, August 24–27.

Pitrebois, Sandra, Michel Denuit, and Jean-François Walhin. 2003b. Setting a Bonus–Malus scale in the presence of other rating factors: Taylor's work revisited. *ASTIN Bulletin: The Journal of the IAA* 33: 419–36. [CrossRef]

Tzougas, George, Spyridon Vrontos, and Nicholas Frangos. 2014. Optimal Bonus–Malus systems using finite mixture models. *ASTIN Bulletin: The Journal of the IAA* 44: 417–44. [CrossRef]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.