

EPiC Series in Computing

Volume 114, 2026, Pages 47–90

Proceedings of the Workshop on Mathematical Modeling and Scientific Computing



Effective Pandemic Management Using Data-Driven Threshold Binning Methods

Nora Granzow[1], Moritz Gubitzer[1], Shiyuan Niu[1], and Tengxiao Gao,[1]
Christina Kuttler[1] and Joachim Draeger[2]

[1] Technical University Munich, Department of Mathematics, Garching, Germany
ge78soy@mytum.de, moritz.gubitzer@tum.de, shiyuan.niu@tum.de, go49boc@mytum.de, kuttler@ma.tum.de
[2] IABG mbH, Ottobrunn, Germany Draeger@iabg.de

Abstract

We discuss the simulation-based generation of epidemics management rules for an SEIRHCD epidemic model with inclusion of a potential non-compliance of the population, and of the economic costs of epidemic countermeasures. A special focus of our research is the necessary categorisation of measurement and input data, which is one of the prerequisites of the intended association rule mining. We consider both an evolutionary search and the application of dynamic programming for finding appropriate bin boundaries. Information criteria help us to prefer a small number of bins. Our results suggest that information-theoretic binning combined with dynamic programming may support the generation of simple and well interpretable rule systems for pandemic countermeasure planning.

1 Introduction

Epidemic management has once again become an important research topic after the COVID-19 pandemic demonstrated how profoundly infectious diseases can affect health systems, economies, and everyday social life on a global scale [33]. The governments reacted with countermeasures such as school closures, contact restrictions, and lockdown. While such measures can help control the spread of disease, they may also impose substantial social and economic costs. For example, a lockdown has been associated with disruptions to employment and economic activity, and restrictive policies can become difficult to sustain if public acceptance decreases over time [4]. Moreover, the effectiveness of countermeasures depends not only on their effect on epidemics but also on behavioural factors such as risk perception, trust, and compliance. Banerjee et al. show that underestimating exponential disease growth is negatively associated with compliance with COVID-19 safety measures [6], while Blair et al. find that during the Ebola epidemic in Liberia lower trust in government was associated with lower compliance with public health interventions [8]. These examples illustrate that epidemic response should not simply consist of activating the strongest possible countermeasure. Instead, intervention strategies should balance health outcomes, economic costs, and compliance-related effects. The aim

V. Turova, A. Kovtanyuk, A. Maslovskaya and J. Zimmer (eds.), MMSC 2026 (EPiC Series in Computing, vol. 114), pp. 47–90

of this study is therefore to support the construction of a decision framework that can reflect these trade-offs and adapt them to different policy preferences. Its rule-based character allows insight into the inner mechanisms.

Compartmental models provide a generic description for the spread of infectious diseases. The SIR model introduced by Kermack and McKendrick divides the population into susceptible, infected, and removed individuals [23]. For the purpose of mitigation management, this basic structure may need to be refined. Policy decisions depend not only on the number of infections, but also on hospital burden, intensive-care demand, and mortality. Accordingly, SEIRHCD-type models extend the classical compartmental framework by including exposed, hospitalised, critical, and deceased compartments [42]. We will follow this approach.

In order to compare different intervention strategies, we simulate a large number of scenarios with varying activation and cancellation thresholds for several countermeasures. During each simulation, a control unit observes the current epidemic state and determines when countermeasures are introduced or released. The resulting trajectories are then evaluated using a total cost function that combines health-related outcomes, economic burden, and non-compliance of the population. In this way, each mitigation strategy is assigned to a cost category, such as low, medium, or high cost. The subsequent rule-mining aims at identifying relationships between threshold choices, countermeasure behaviour, and the resulting cost category.

An association rule mining requires categorical input data, whereas e.g. the activation and cancellation thresholds are continuous values. These thresholds must therefore be discretised before rules can be extracted. If the corresponding bin boundaries are chosen arbitrarily, potential patterns in the simulation data concerning input-output-relationships may be lost. Instead of relying only on fixed, equally spaced intervals, this study uses thus a data-driven binning approach based on mutual information between threshold bins and the cost category. Since the threshold values are ordered and the problem is decomposable into smaller subproblems — i.e. situations in which a smaller number of suitable bin intervals have to be found — the optimal bin boundaries for a given number of bins can be found exactly using dynamic programming [7]. The final number of bins is then selected using the Bayesian Information Criterion [35], so that the resulting discretisation balances cost-relevant information with interpretability and avoids unnecessarily complex binning.

The paper is structured as follows. Chapter 2 reviews related work on epidemic modelling and rule-based decision support. Chapter 3 introduces the existing SEIRHCD model framework and the total cost function. Chapter 4 describes the data processing pipeline. The data-driven threshold binning approaches based on mutual information are discussed in Chapter 5. We first discuss an evolutionary binning approach and then show how the ordered structure of the threshold values allows the binning problem to be solved exactly by dynamic programming, with the final number of bins selected using the Bayesian Information Criterion. Chapter 6 describes how the binned simulation data are used for association rule mining and policy extraction. Finally, Chapter 7 summarises the main findings, limitations, and possible directions for future work.

2 Literature review

Epidemic models have been widely used to evaluate and compare different intervention strategies. Dhaoui et al. evaluated the effects of non-pharmaceutical interventions using an extended compartmental model. Similarly, Thompson et al. examined the effects of different intervention strategies on epidemic outcomes [14,41]. Other studies considered enriched assessments of outcomes when evaluating intervention strategies. Perkins et al. considered both epidemic con-

trol and the wider costs of interventions, while Salehi et al. included public health, healthcare capacity, and economic effects in their evaluation [27, 34]. Ramponi et al. further combined epidemiological and economic indicators to study the trade-off between health and economic outcomes [31]. For assessing the effects of epidemic countermeasures from the perspective of risk mitigation measures, [15, 17] may serve as a theoretical framework.

Public compliance has also been included in epidemic models because the effectiveness of an intervention depends partly on whether people follow it. Parkinson et al. analysed the role of non-compliance in determining the effects of intervention measures [26]. Compliance aspects were also analysed in [18], which is available in a comprehensive report version as well [19]. An extension of the psychological perspective from subconscious processes to cognitive aspects is given in [16]. Though [16] relates to social engineering as domain of application, the presented concept may be easily adaptable to an epidemiological context.

For discretisation strategies in association rule mining, different approaches have been used to handle continuous variables. Equal-width and equal-frequency binning are among the most common approaches. Putri et al. compared these methods on several datasets, including COVID-19 and dengue fever data, and showed that the choice of discretisation method can affect subsequent classification results [28]. In association rule mining, He et al. also compared equal-width and equal-frequency discretisation and found that the choice of method can affect the generated rules [21]. Other studies have used supervised discretisation methods that include class information when selecting cut points. These approaches often use measures such as Information Gain, Entropy, or Mutual Information. For example, Trung et al. proposed a non-parametric discretisation method based on Dynamic Programming to globally optimise cut-point selection [29].

In addition, evolutionary optimisation methods such as Genetic Algorithms have been proposed to search for optimal cut points for continuous variables. These methods are capable of performing global optimisation within complex search spaces and have therefore been increasingly applied in classification, data mining, and knowledge discovery tasks [30]. For example, Ramírez-Gallego et al. employed a Genetic Algorithm to optimise cut points for multiple continuous variables simultaneously. Their experimental results demonstrated that the proposed method produced higher-quality discretisation and further improved classification performance [30]. These studies show that different optimisation methods can be used to improve the discretisation of continuous variables.

Beyond discretisation, previous research has also explored how simulation and optimisation methods can support the selection of epidemic intervention strategies. For example, Padmanabhan et al. proposed a reinforcement learning-based decision support system that dynamically adjusts interventions according to the current epidemic state [25]. Similarly, Reymond et al. used multi-objective reinforcement learning to balance infections and hospitalisations against the social burden caused by intervention measures [32]. However, the strategies generated by these methods can be complex, making the decision-making processes for policymakers difficult to understand. Additionally, though neural networks may generate highly performing solutions, they typically miss insight into the mechanisms of the generated solutions. Therefore, some studies have focused more on the interpretability of pandemic policies. Custode and Iacca combined reinforcement learning with evolutionary computation to generate simple and interpretable intervention strategies from epidemic simulations. Their work shows that pandemic decision support should not only identify effective intervention strategies but also make these strategies understandable to policymakers [12].

Association Rule Mining (ARM) was first introduced by Agrawal et al. to discover relationships among variables in large datasets [1]. ARM can generate understandable decision rules in

a straightforward manner. It has been widely applied in healthcare, bioinformatics, and clinical decision support [3]. ARM has also been used in COVID-19 and other infectious disease research to discover interpretable relationships in complex epidemic data. Tandan et al. applied association rule mining to COVID-19 patient data and extracted rules between symptoms, age, gender, chronic diseases, and mortality [40]. Somyanonthanakul et al. also used ARM to analyse symptoms, age, and underlying diseases of COVID-19 patients, and further used the identified important factors for epidemic forecasting and healthcare resource management [38]. Sinisterra-Sierra et al. used a multi-objective evolutionary algorithm to extract association rules and causal association rules from large-scale COVID-19 data. Since the dataset contained continuous variables, they applied quintile-based discretisation to divide the numerical data into five intervals with approximately equal support [37].

Overall, ARM has proven useful for extracting intuitive and understandable relationships from complex epidemic data. In epidemic management, rule-based approaches can extract relationships between intervention strategies and the resulting total costs from large numbers of simulation outcomes, providing decision-makers with intuitive and interpretable policy recommendations. In addition, the quality of the generated rules strongly depends on the discretisation of the input variables. Therefore, optimising the discretisation process has become an important research direction for improving rule quality [39]. In some way, this study contributes to this research. It uses simulation data generated by the SEIRHCD model and discretises the activation and cancellation thresholds of the countermeasures based on different binning strategies. Their effects on the quality of the generated rules are analysed.

3 Model framework

This section introduces the simulation model providing the epidemic-management framework for the present study. The framework combines an SEIRHCD compartmental model with a countermeasure control unit, behavioural non-compliance of the population, and a total cost evaluation. The later rule-mining procedure will operate on simulation outputs with activation thresholds, cancellation thresholds, countermeasure durations, and cost labels.

The present work adopts the modelling framework developed in previous case-studies [5, 13, 43] and focuses on improving the preprocessing steps before association rule mining, especially the discretisation of continuous threshold values. The following subsections summarise the model components that are necessary for understanding the simulation data and the dynamic threshold binning approach introduced later.

3.1 The SEIRHCD model

The epidemiological submodel of the simulation framework is an extension of the classical SEIR model. In addition to the susceptible (S), exposed (E), infected (I), and recovered (R) compartments, the model includes hospitalised individuals (H), people in critical condition (C), and deceased individuals (D). This results in an SEIRHCD model that is well suited for analysing pandemic management strategies.

Figure 1 summarises the epidemic transitions. Susceptible individuals become exposed through contact with infectious individuals; exposed individuals subsequently become infectious. Infectious individuals either recover or require hospital treatment. Hospitalised individuals may recover or enter critical care; critical patients either return to hospital care or die. Recovered individuals may become susceptible again through loss of immunity. Vaccination is

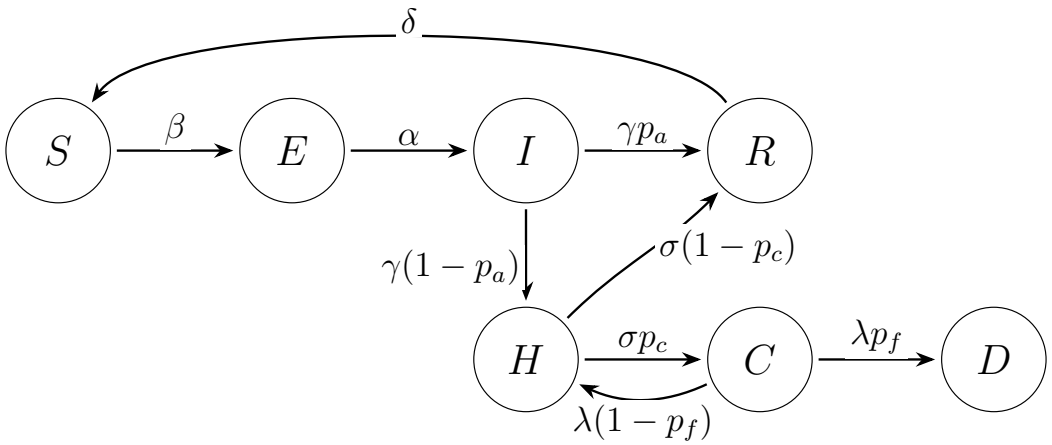


Figure 1: Flowchart of the epidemic model with compartments susceptible S, exposed E, infected I, recovered R, hospitalised H, critical condition C, and deceased D. See Table 1 for a description of the parameters α, β, γ, δ, σ, λ, p_a, p_c, and p_f.

disabled to represent an outbreak setting without an available vaccine, focusing on rules for activating and lifting non-pharmaceutical interventions. The corresponding rates and branching probabilities are specified in Table 1.

Definition 3.1 (ODE system for the SEIRHCD model)**.** *The system of ordinary differential equations governing the SEIRHCD model is given by:*

$$\begin{aligned}
\frac{dS(t)}{dt} &= -\beta S(t)I(t) + \delta R(t),\\
\frac{dE(t)}{dt} &= \beta S(t)I(t) - \alpha E(t),\\
\frac{dI(t)}{dt} &= \alpha E(t) - \gamma I(t),\\
\frac{dH(t)}{dt} &= \gamma(1-p_a)I(t) + \lambda(1-p_f)C(t) - \sigma H(t),\\
\frac{dC(t)}{dt} &= \sigma p_c H(t) - \lambda C(t),\\
\frac{dR(t)}{dt} &= \gamma p_a I(t) + \sigma(1-p_c)H(t) - \delta R(t),\\
\frac{dD(t)}{dt} &= \lambda p_f C(t).
\end{aligned} \tag{1}$$

Births and natural deaths are neglected. Throughout this study, the compartments $S, E, I, H,\ C,\ R, D$ denote dimensionless proportions of the initial population, and time is measured in days. Disease-related deaths are tracked cumulatively in D as a proportion of the initial population. The total tracked population proportion, including deceased individuals, is

$$N(t) = S(t) + E(t) + I(t) + H(t) + C(t) + R(t) + D(t).$$

Summing the seven compartment equations (1) yields

$$\frac{dN(t)}{dt} = 0.$$

The normalised initial conditions satisfy $N(0) = 1$; hence $N(t) = 1$ throughout the simulation.

The epidemiological parameters used in the framework and their numerical values are summarised in Table 1. The baseline transmission rate β is varied between simulation experiments to represent different levels of epidemic severity and is therefore not assigned a single fixed value in the table.

Table 1: Overview of the fixed parameters used in the SEIRHCD model, including their values, units, and descriptions.

Parameter	**Value**	**Unit**	**Meaning**
$\alpha \in \mathbb{R}^+$	0.3	days^{-1}	Rate of becoming infectious, i.e. incubation rate.
$\beta \in \mathbb{R}^+$	–	days^{-1}	Transmission rate, i.e. infection rate.
$\gamma \in \mathbb{R}^+$	0.071	days^{-1}	Rate of leaving the state "infectious".
$\delta \in \mathbb{R}^+$	0.0055	days^{-1}	Rate of becoming "susceptible" again after recovery.
$\sigma \in \mathbb{R}^+$	0.5	days^{-1}	Rate of leaving the state "hospitalised", either by recovery or transfer to the ICU.
$\lambda \in \mathbb{R}^+$	0.5	days^{-1}	Rate of leaving the state "critical condition", either by becoming hospitalised again or by dying.
$p_a \in [0, 1]$	0.79	–	Fraction of infected individuals needing no hospital treatment.
$p_c \in [0, 1]$	0.12	–	Fraction of hospitalised individuals switching to a critical state.
$p_f \in [0, 1]$	0.33	–	Fraction of patients in critical condition who ultimately die.

3.2 Control unit

For managing countermeasures against the epidemic, a control unit is introduced. The control unit assesses the situation using a scalar score. Countermeasures are activated once this score crosses predefined activation thresholds and are cancelled once it falls below corresponding cancellation thresholds.

Definition 3.2 (control score)**.** *The control unit uses the score*

$$\Phi_{\text{CU}}(t) = \frac{\epsilon w_I I(t) + w_H H(t) + w_C C(t) + w_D D(t)}{w_I + w_H + w_C + w_D}. \tag{2}$$

Here $\epsilon \in [0, 1]$ models imperfect observation of infections, and w_I, w_H, w_C, w_D are non-negative weights.

The control score $\Phi_{\text{CU}}(t)$ is dimensionless: the compartment proportions, observation factor ϵ, and weights are dimensionless. Activation thresholds T_i and cancellation thresholds $\widetilde{T}_i$ are values of this score and do not directly represent percentages of infectious individuals. The control score weights are chosen as

$$(w_I, w_H, w_C, w_D) = (0.75, 0.10, 0.10, 0.05), \qquad \epsilon = 0.6.$$

Thus, the score is typically dominated by the infected compartment, while hospitalised, critical, and deceased individuals still contribute to the assessment of the epidemic situation. The associated thresholds and effectiveness coefficients for these interventions are defined as follows.

Definition 3.3 (Thresholds, Efficiency)**.** *When the score function $\Phi(t)$ exceeds thresholds*

$$0 < T_1 < T_2 < T_3 \leq T_{\max}, \tag{3}$$

countermeasures will be set into effect, i.e. the transmission rate β will be reduced by multiplication with the so-called efficiency coefficients

$$0 < \eta_3 < \eta_2 < \eta_1 < 1.$$

When the population is fully compliant, and the i-th countermeasure is in effect, the effective transmission rate will be

$$\eta_i \beta < \beta.$$

Smaller η corresponds to harsher countermeasures. Countermeasures may also be released. The thresholds to release the countermeasures are pairwise smaller than the thresholds to activate them, as introduced in (3),

$$0 \leq \widetilde{T}_1 < \widetilde{T}_2 < \widetilde{T}_3 < T_{\max} \quad \textit{and} \quad \widetilde{T}_i < T_i.$$

We denote by

$$\eta(t) \in \{1, \eta_1, \eta_2, \eta_3\}$$

the efficiency level in effect at time t, where $\eta(t) = 1$ corresponds to no active countermeasure.

In the simulations, the choice of the sampling range $T_{\max} = 0.3$ is discussed in Section 4.1 and three predefined values of η are used, corresponding to three increasingly strict countermeasures. [5]

$$\eta_1 = 0.85, \qquad \eta_2 = 0.45, \qquad \eta_3 = 0.15.$$

3.3 Non-compliance sub-model

In practice, the effectiveness of countermeasures may be limited by non-compliance within the population. Even if a countermeasure is formally introduced, its epidemiological effect depends on the extent to which individuals follow the imposed restrictions. The model therefore includes a time-dependent non-compliance factor

$$A(t) \in [0, 1],$$

where small values of $A(t)$ correspond to high compliance and large values correspond to reduced compliance. This component is intended to capture behavioural reactions to the relationship between the current epidemic situation and the strictness of the implemented countermeasures. In the inherited framework, non-compliance is driven by the population's perception of whether the current restrictions are appropriate. The implemented simulation code represents this perception by a population score φ_{pop}, which is computed from the current epidemic state and the transition rates between epidemiological compartments. In contrast to the control unit score in Section 3.2, the population score does not use the compartment levels I, H, C, D directly. Instead, it uses quantities related to new infections, recoveries, critical cases, and deaths.

Definition 3.4 (Population score)**.** *The dimensionless population perception score is defined by*

$$\varphi_{\text{pop}}(t) = t_{\text{ref}} \frac{\epsilon \kappa_I \alpha E(t) + \kappa_H \gamma p_a I(t) + \kappa_C \sigma p_c H(t) + \kappa_D \lambda p_f C(t)}{\kappa_I + \kappa_H + \kappa_C + \kappa_D}. \tag{4}$$

Here, $t_{\mathrm{ref}} = 1\,\mathrm{day}$ *is a fixed reference time, and* $\epsilon = 0.6$ *is a dimensionless factor modelling imperfect observation of infections. The weights* $\kappa_I, \kappa_H, \kappa_C, \kappa_D$ *are dimensionless and non-negative, with positive sum, and describe the relative importance assigned by the population to different aspects of the epidemic situation. Since the weighted transition-rate expression has units* day^{-1}*, multiplication by* t_{ref} *makes the score dimensionless.*

The dimensionless population perception score $\varphi_{\mathrm{pop}}(t)$ is compared with the dimensionless restriction level $1 - \eta(t)$. Smaller values of $\eta(t)$ correspond to stricter countermeasures. If the restrictiveness $1 - \eta(t)$ is larger than the population score, the measures are perceived as too strict relative to the current situation. This perceived over-restrictiveness leads to a tendency not to comply with the rules. For the simulations, this effect is represented by two state variables, $A(t)$ and $B(t)$. The auxiliary variable $B(t)$ stores an exponentially decaying memory of past perceived over-restrictiveness. It satisfies

$$\dot{B}(t) = -bB(t) + b\max\{0, 1 - \eta(t) - \varphi_{\mathrm{pop}}(t)\}, \tag{5}$$

where $B(t)$ is dimensionless and $b > 0$ has units day^{-1}, controlling how quickly past events are forgotten. Recent mismatches between restriction stringency and perceived epidemic severity therefore have a stronger effect than events further in the past. The non-compliance factor $A(t)$ then adapts towards the memory variable according to

$$\dot{A}(t) = M\big(mB(t) - A(t)\big). \tag{6}$$

The parameter $M > 0$ has units day^{-1} and controls the adaptation rate of the dimensionless non-compliance factor $A(t)$. The dimensionless parameter $m > 0$ scales the influence of the memory variable. In the numerical experiments used in this study, we use the same non-compliance parameter values as in [5]:

$$m = 1, \qquad b = 0.1\,\mathrm{day}^{-1}, \qquad M = 0.2\,\mathrm{day}^{-1}, \qquad A(0) = A_0 = 0.$$

Thus, the population is assumed to be fully compliant at the beginning of the simulation, and non-compliance only develops in response to countermeasures that are perceived as too strict relative to the epidemic situation. Non-compliance also affects the transmission dynamics. In the model, the effective transmission rate combines the formal countermeasure efficiency $\eta(t)$ with the current non-compliance level $A(t)$. It is given by

$$\beta_{\mathrm{eff}}(t) = \beta\left[\eta(t) + \big(1 - \eta(t)\big)A(t)\right]. \tag{7}$$

Thus, if $A(t) = 0$, the countermeasure has its full intended effect and $\beta_{\mathrm{eff}}(t) = \beta\eta(t)$. If $A(t) = 1$, the countermeasure is completely ineffective and $\beta_{\mathrm{eff}}(t) = \beta$. The formulation above follows the implemented simulation pipeline used for the present study. The implementation uses time measured in days and numerical rate values expressed per day. The fixed reference time $t_{\mathrm{ref}} = 1\,\mathrm{day}$ therefore has numerical value one. Making this scaling explicit preserves the numerical population scores and does not require changes to the simulation results. For the rule-mining dataset, however, the relevant quantities are those actually generated by the simulation code. We therefore use the Equations (4)–(7) as the model description for this framework. Beyond potential non-compliance, the model also includes a second non-health related component: the economic cost arising from both the pandemic itself and the implemented countermeasures.

3.4 Economic sub-model

The economic sub-model assigns a cost to the combined effect of restrictions and illness-related workforce reduction. An epidemic countermeasure may reduce the number of new infections, but it usually also reduces economic activity. Conversely, weak restrictions may preserve economic activity in the short run, but can allow more actual infections, which also reduces the available workforce. The model evaluates this trade-off through a cumulative economic cost.

Definition 3.5 (Economic cost)**.** *Let*

$$W(t) = S(t) + R(t)$$

denote the proportion of the initial population that is either susceptible or recovered. This dimensionless quantity is used as a proxy for the available workforce. The instantaneous economic cost is defined as

$$c_E(t) = \rho\, f_\eta(\eta(t))\, f_W(W(t)),$$

where f_η measures the economic cost of restrictions and f_W measures the cost caused by workforce loss. In the simulation model, $\rho = 0.16$ [5], and

$$f_\eta(\eta) = \begin{cases} \left(1 + \exp\left[-\log(\eta^{-1} - 1) - 0.1\right]\right)^{-1}, & 0 < \eta < 1, \\ 0, & \eta = 1. \end{cases}$$

$$f_W(W) = \left(1 + \exp\left[-\log(W^{-1} - 1) - 0.3\right]\right)^{-1}.$$

Here, smaller values of η correspond to stricter countermeasures and therefore larger restriction costs, while smaller values of W correspond to a larger loss of available workforce. In continuous notation, the cumulative economic cost can be written as

$$C_E^{\text{cont}} = \int_0^{t_f} c_E(t)\, dt.$$

Numerically, the economic cost is approximated as discrete sum

$$C_E = \sum_{k=1}^{K} c_E(t_k),$$

over the simulated time points $t_1, \ldots, t_K$. This discrete value is the economic-cost quantity used later in the total cost function and in the rule-mining dataset.

3.5 Total cost function

Each simulated scenario is assigned one total cost value. We adopt the structure of the total cost function from the previous case study [5], using its reduced formulation for the epidemiological component. The original formulation considered infections (I), hospitalisations (H), critical cases (C), and deaths (D) as separate epidemiological cost terms. However, their analysis showed that H, C, and D are strongly correlated with I and therefore add little independent information to the cost evaluation. For this reason, we simplify the cost function and use the infectious population proportion $I(t)$ to represent the epidemiological burden. The total cost combines three components: infection burden (see above), non-compliance, and economic loss.

Let s denote a simulated scenario and let $t_1, \ldots, t_K$ be the sampled simulation time points. The raw infection and non-compliance costs are defined as threshold-excess costs,

$$R_I^s = \sum_{k=1}^{K} \max\{0, I^s(t_k) - \tau_I\}, \qquad R_A^s = \sum_{k=1}^{K} \max\{0, A^s(t_k) - \tau_A\}.$$

Here, τ_I and τ_A are dimensionless cost thresholds for the infectious population proportion and the non-compliance level, respectively. Values below the corresponding threshold do not contribute to the cost, while values above the threshold are penalised according to both their magnitude and duration. The economic component is treated differently because the model already returns an aggregated economic cost value. Therefore, no additional threshold-excess calculation is applied to the economic component. The raw components R_I^s, R_A^s, and R_E^s are later normalised. The normalised versions of the raw infection, non-compliance, and economic components are denoted as C_I^s, C_A^s, and C_E^s. They are combined into a weighted total cost of the form

$$C_{\text{total}}^s = \omega_I C_I^s + \omega_A C_A^s + \omega_E C_E^s, \qquad \omega_I + \omega_A + \omega_E = 1, \qquad \omega_I, \omega_A, \omega_E \geq 0.$$

The weights describe the relative importance assigned to infection control, societal acceptance, and economic burden. They can therefore be interpreted as policy preferences, allowing decision makers the option to choose where to place more emphasis. The chosen values of the thresholds τ_I, τ_A, of the weights $\omega_I, \omega_A, \omega_E$, the adopted normalisation procedure, and the resulting cost categories are discussed in the next section.

4 Adapted simulation and cost label construction

Before applying threshold binning and association rule mining, the continuous simulation output is converted into a labelled rule-mining dataset. This section describes the corresponding preprocessing pipeline, which prepares the generated simulation results for the dynamic-programming threshold binning introduced in the next section. The main improvements compared with [5] are the recalibrated countermeasure-threshold sampling range, the reproducible choice of the infection and non-compliance thresholds used in the cost function, and the automated construction of low-, medium-, and high-cost labels. Unless stated otherwise, all quantities in the following sections are computed separately for each transmission setting

$$\beta \in \{0.20, 0.35, 0.50, 0.80\}.$$

4.1 Recalibrated threshold sampling range

Recall from Section 3.2 that the countermeasure level $i \in \{1, 2, 3\}$ is activated when the control score $\Phi_{\text{CU}}(t)$ exceeds its activation threshold T_i, and it is released when the score falls below the corresponding cancellation threshold $\widetilde{T}_i$. We follow the choices made in [5]. Thus, in the following we will use the values

$$(w_I, w_H, w_C, w_D) = (0.75, 0.1, 0.1, 0.05) \quad \text{and} \quad \epsilon = 0.6.$$

The weights sum to one and the score becomes

$$\Phi_{\text{CU}}(t) = 0.45\, I(t) + 0.10\, H(t) + 0.10\, C(t) + 0.05\, D(t).$$

For each simulation scenario, the six threshold values

$$\{\widetilde{T}_1, T_1, \widetilde{T}_2, T_2, \widetilde{T}_3, T_3\}$$

are sampled independently from a uniform distribution on $(0, S)$ and then sorted. Thereafter, the values are assigned as

$$\widetilde{T}_1 = x_{(1)}, \qquad T_1 = x_{(2)}, \qquad \widetilde{T}_2 = x_{(3)}, \qquad T_2 = x_{(4)}, \qquad \widetilde{T}_3 = x_{(5)}, \qquad T_3 = x_{(6)}.$$

This guarantees w.r.o.g. the escalation ordering

$$\widetilde{T}_1 < T_1 < \widetilde{T}_2 < T_2 < \widetilde{T}_3 < T_3,$$

so that every cancellation threshold lies below its corresponding activation threshold. Since T_1, T_2, T_3 are the second, fourth, and sixth order statistics of six uniform draws, their expected values are

$$\mathbb{E}[T_1] = \frac{2}{7}S, \qquad \mathbb{E}[T_2] = \frac{4}{7}S, \qquad \mathbb{E}[T_3] = \frac{6}{7}S. \tag{8}$$

The framework in [5] used $T_{max} = 0.5$. However, Equation (2) shows that the reachable control-score range is much smaller than the full interval $(0, 0.5)$. To estimate this range from the simulated data, we evaluated the score formula using the maximum observed values of I, H, C, and D in each trajectory. The largest resulting values for each β are:

β	0.20	0.35	0.50	0.80
Maximum observed score	0.107	0.175	0.209	0.243

Equation (8) gives

$$\mathbb{E}[T_1] \approx 0.14, \qquad \mathbb{E}[T_2] \approx 0.29, \qquad \mathbb{E}[T_3] \approx 0.43.$$

The expected thresholds for the second and third countermeasures therefore lie above the observed score ceiling for all considered transmission rates in the previous setup. This explains why the third countermeasure was almost never activated: the three-countermeasure system was effectively reduced to a smaller set of usable policy levers. To assure an adequate escalation structure, the threshold sampling range is thus set to

$$T_{max} = 0.3.$$

Then

$$\mathbb{E}[T_1] \approx 0.09, \qquad \mathbb{E}[T_2] \approx 0.17, \qquad \mathbb{E}[T_3] \approx 0.26.$$

This places the first and second countermeasures inside the empirically relevant score range, while the third countermeasure remains primarily a severe-scenario intervention. Under this recalibration, countermeasure three (CM3) is still inactive in the mild setting but becomes reachable for higher transmission rates. Figure 2 compares the activation frequencies under the old and recalibrated sampling ranges.

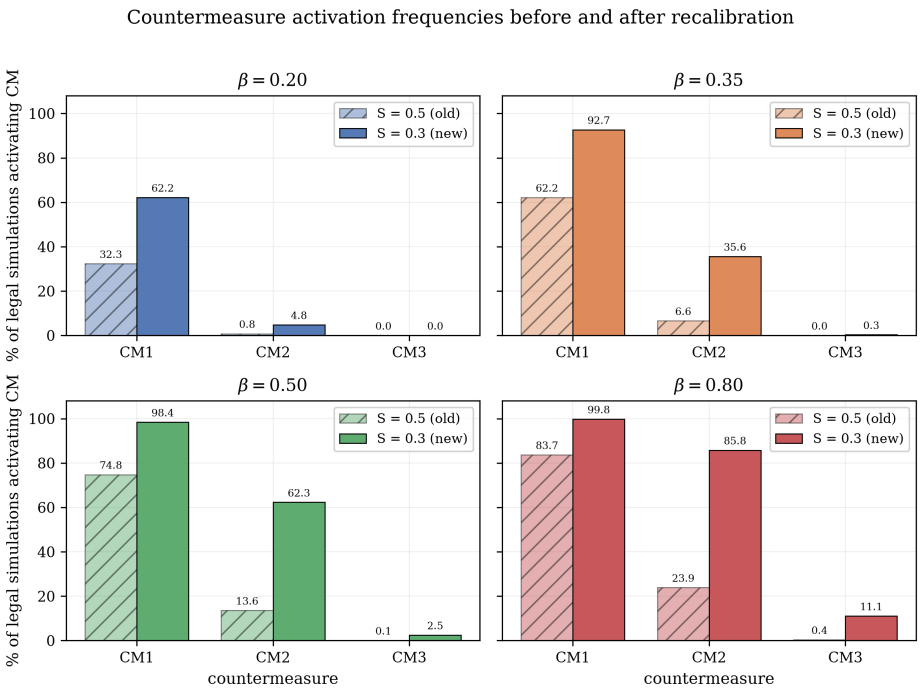


Figure 2: Percentage of selected simulations in which each countermeasure is activated, comparing the old sampling range $S = 0.5$ and the recalibrated sampling range $S = 0.3$ (bars are cumulative activation rates).The endpoint S corresponds to text's T_{max}.

The recalibration also changes the non-compliance dynamics, since non-compliance reacts to active restrictions. This is a desired consequence. Under the old range, parts of the countermeasure system were too rarely used for the behavioural and rule-mining components to reflect the intended three-level control structure.

4.2 Simulation output and scenario selection

For each transmission setting, we have generated

$$N = 10{,}000$$

scenarios. Scenario s is defined by the sorted threshold vector

$$\left(\widetilde{T}_1^s, T_1^s, \widetilde{T}_2^s, T_2^s, \widetilde{T}_3^s, T_3^s\right),$$

together with the corresponding trajectory over $[0, T_{\text{end}}]$, where $T_{\text{end}} = 1000\,\text{days}$. For the infection and non-compliance cost calculations, the trajectories are sampled at $K = 1000$ equally spaced time points. The relevant trajectory variables for cost construction are the infectious population proportion $I^s(t)$ and the dimensionless non-compliance level $A^s(t)$. The simulation also stores the levels of the other epidemic compartments, the cumulative economic cost value $E^s \geq 0$, and the set of activation and cancellation events occurring in scenario s.

Scenarios with missing economic cost values are first excluded from cost construction. Among the remaining scenarios, let a_i^s and c_i^s denote the numbers of activation and cancellation events, respectively, for countermeasure i in scenario s within the simulation horizon. The selected scenario set is

$$\mathcal{L} = \{s : a_i^s = c_i^s \quad \text{for all } i \in \{1,2,3\}\}, \qquad n = |\mathcal{L}|.$$

Thus, every activation must have a corresponding cancellation within the observation period. If any activation remains unmatched, the entire scenario is excluded. A countermeasure that never activates satisfies the criterion automatically, with $a_i^s = c_i^s = 0$.

The intended application is a policy rulebook supporting decisions about which countermeasure level to activate, when to activate it, and when to lift it, while balancing epidemiological, economic, and non-compliance costs. Following the previous framework [5], we therefore focus on temporary intervention histories in which both activation and subsequent release are observed. An intervention history with no observed release, such as a lockdown still active at the simulation horizon, provides an incomplete basis for this intended guidance, even if its cost over the observation period is favourable.

This selection is a restriction of the study's scope, rather than a requirement for mathematical or numerical validity. A measure still active at the horizon may be lifted later or may remain necessary because of ongoing epidemic activity. Conversely, release of all measures does not imply that the epidemic has ended. All subsequent cost normalisations, cost-label construction, threshold binning, and rule mining use the selected set $\mathcal{L}$, unless stated otherwise.

The resulting associations apply to this selected set; their sensitivity to including scenarios with ongoing interventions has not been assessed.

4.3 Cost thresholds τ_I and τ_A

The total cost function contains dimensionless thresholds τ_I and τ_A for the infectious population proportion and the non-compliance level, respectively. These thresholds determine from which point onward infection burden and non-compliance contribute to the scenario cost. For each selected simulation s, define

$$I_{\max}^{(s)} = \max_t I^{(s)}(t), \qquad A_{\max}^{(s)} = \max_t A^{(s)}(t).$$

The cost thresholds τ_I and τ_A apply directly to the infectious population proportion $I(t)$ and the dimensionless non-compliance factor $A(t)$, respectively. Activation and cancellation thresholds instead apply to the weighted control score $\Phi_{\mathrm{CU}}(t)$.

In [5], the cost thresholds τ_I and τ_A were chosen by visual inspection of the empirical distributions. Here, the thresholds are formally derived based on the empirical statistical distributions of $I_{\max}$ and $A_{\max}$ after their recalibration. These calculations are made before the cost weights are applied for computing the total cost. The purpose of this step is to define interpretable tolerance levels for infection burden and non-compliance in a consistent way across transmission settings. The resulting choices are shown in Table 2. The non-compliance threshold is kept fixed at $\tau_A = 0.20$ for all transmission settings, while the infection threshold is transmission-specific.

Table 2: The cost-function thresholds used for the assessment. The infection thresholds are transmission-specific tolerance levels derived from the $I_{\max}$ distributions. The non-compliance threshold is fixed at $\tau_A = 0.20$.

β	τ_I	Share $I_{\max} > \tau_I$	τ_A	Share $A_{\max} > \tau_A$
0.20	0.20	88.7%	0.20	4.8%
0.35	0.30	90.4%	0.20	35.0%
0.50	0.375	85.0%	0.20	61.1%
0.80	0.45	92.0%	0.20	84.6%

The exceedance shares show that τ_I does not only isolate rare extreme simulations but it also penalises sustained infection burden above a transmission-specific tolerance level. The non-compliance threshold behaves differently. Whereas only a small fraction of mild simulations exceed τ_A, non-compliance above τ_A becomes common in severe settings. This seems to be natural, because stronger and more frequent countermeasures lead to stronger non-compliance responses.

We keep $\tau_A = 0.20$ fixed across transmission settings. In this way, it represents the same behaviour in all scenarios: non-compliance below this value is treated as manageable background variation, while values above it indicate a substantial loss of adherence. Unlike τ_I, which must be transmission-specific because the infection burden naturally increases with β, $A(t)$ is already measured on a common $[0, 1]$ scale. Keeping τ_A fixed therefore makes the non-compliance component comparable across transmission settings.

4.4 Normalisation of total cost function components

Each simulated scenario must be assigned one total cost value before it can be used for rule mining as defined in 3.5. For every selected scenario $s \in \mathcal{L}$, the infection and non-compliance components are computed as threshold-excess costs. This means that values below a chosen threshold do not contribute to the cost, while values above the threshold are accumulated over time:

$$R_I^s = \sum_{t=1}^{K} \max\{0, I^s(t) - \tau_I\}, \qquad R_A^s = \sum_{t=1}^{K} \max\{0, A^s(t) - \tau_A\}. \tag{9}$$

The economic component already returns an aggregated economic cost value. Therefore, no additional threshold or time integration is applied. The three raw components R_I^s, R_A^s in (9), and R_E^s are not directly comparable, because they have different numerical scales. For example, the economic cost may be much larger in absolute value than the infection or non-compliance cost, without necessarily being more important. Therefore, each component is normalised before the total cost is formed. For this purpose, we divide by a robust reference value. For each component $X \in \{I, A, E\}$, this reference value is the 99th percentile of the raw component costs over all selected simulations:

$$r_X = Q_{0.99}\left(\{R_X^s : s \in \mathcal{L}\}\right).$$

If this value is zero, the empirical maximum is used as a fallback. The normalised component is then

$$C_X^s = \min\left\{\frac{R_X^s}{r_X}, 1\right\}.$$

This gives $C_X^s \in [0,1]$. Values above the 99th percentile are clipped to one, so that a small number of extreme simulations cannot stretch the scale for all other scenarios. The total cost of scenario s is finally defined as a weighted sum of the three normalised components:

$$C_{\text{total}}^s = \omega_I C_I^s + \omega_A C_A^s + \omega_E C_E^s.$$

In the baseline setting, all three components are weighted equally:

$$\omega_I = \omega_A = \omega_E = \frac{1}{3}.$$

The weights can be interpreted as policy preferences. Increasing ω_I gives more importance to reducing infections, increasing ω_A gives more importance to avoiding non-compliance, and increasing ω_E gives more importance to limiting economic loss.

4.5 Thresholds for the total cost function

After computing the continuous total cost C_{total}^s, each selected scenario must be assigned to a categorical cost label. This is necessary because a discrete target variable is mandatory for the rule-mining. We therefore introduce two cost thresholds,

$$c_\ell < c_h,$$

which split the total cost values into three classes:

$$\text{Bin}^s = \begin{cases} \text{low}, & C_{\text{total}}^s \le c_\ell, \\ \text{med}, & c_\ell < C_{\text{total}}^s \le c_h, \\ \text{high}, & C_{\text{total}}^s > c_h. \end{cases}$$

These thresholds were selected with the help of an automated search over candidate cut points. The candidates are based on empirical percentiles of the total cost distribution. Among these candidates, we choose thresholds that satisfy three conditions: first, each bin must contain a sufficient number of scenarios, so that the later rule-mining step does not rely on very small classes. Second, the bins must be ordered by their mean total cost, meaning that the low-cost bin should have a smaller average cost than the medium-cost bin, and the medium-cost bin should have a smaller average cost than the high-cost bin. Third, the most common cost region is kept in the medium bin. This prevents the typical simulation outcome from being classified as either unusually good or unusually bad. Among all feasible threshold pairs, the selected pair is thus the one that best balances between these three requirements. The resulting thresholds are shown in Table 3.

Table 3: Optimised cost-bin thresholds and resulting class counts (over selected simulations).

β	low|medium	medium|high	low	medium	high	selected n
0.20	0.231	0.333	3474	4376	421	8271
0.35	0.333	0.442	2795	3308	1251	7354
0.50	0.336	0.422	2814	3019	1030	6863
0.80	0.364	0.453	1811	3371	1062	6244

The resulting thresholds are sensible for the following reason. The low-, medium-, and high-cost labels remain clearly ordered by total cost, while all three labels are still represented in

the data. In particular, the high-cost class is not defined only as a tiny extreme tail, and the medium class represents the typical cost region. This gives the dynamic-programming binning step a stable and meaningful target variable: it learns threshold intervals that are informative about whether a scenario leads to low, medium, or high total cost.

Table 4: Parameter values passed to binning.

Symbol	**Meaning**	**Value**
T_{max}	threshold sampling interval	$[0, 0.3]$
N, K	number of scenarios and time steps	10,000, 1000
τ_A	non-compliance threshold	0.20 for all β
τ_I	infection threshold	0.20/0.30/0.375/0.45 for $\beta = 0.20/0.35/0.50/0.80$
$\omega_A, \omega_I, \omega_E$	baseline cost weights	$\frac{1}{3}, \frac{1}{3}, \frac{1}{3}$
$Q_{0.99}$	normalisation reference	per-component, per-β 99th percentile, clipped
$0.05, 0.80$	cost-bin support bounds	minimum and maximum allowed class share
$0.60, 0.30, 0.10$	cost-bin score weights	variance explained, entropy balance, separation

The bin boundaries and the association rules are both derived from the same simulation set. Indeed, the aim of this study is to assess the *quality of the rules themselves* – how reliably the extracted rules describe the structure of the simulated input–output relationship – rather than to predict the cost category of unseen scenarios. Both have the same exploratory purpose: to generate an interpretable rule system from the simulation data. Correspondingly, the values reported in Section 6 should be read as descriptive measures of this in-sample structure, not as out-of-sample predictive performance.

5 Threshold binning

The ultimate goal of this work is to derive interpretable decision rules for pandemic management using association rule mining. However, algorithms such as `apriori()` work on categorical data. The simulation framework produces continuous activation and cancellation thresholds for each countermeasure as well as a continuous total cost value. These quantities therefore have to be discretised before rule mining can be applied. After discretisation, rule-mining can identify associations of the form

$$\{\text{CM1 activation threshold in bin 2}\} \Rightarrow \{\text{total cost is low}\}.$$

In [5], the threshold values for countermeasure activation and cancellation were divided into five equally spaced bins. This provides a simple baseline and makes the resulting rules easy to interpret. However, equal-width binning ignores the actual structure of the simulation data. In that case, potentially useful information is already lost before the rule-mining step starts. A second limitation is that activation and cancellation thresholds were binned in the same way. While this is convenient, the two quantities describe different decisions. Activation thresholds determine when countermeasures are introduced, whereas cancellation thresholds determine when they are released again. Since cancellation thresholds are necessarily smaller than the corresponding activation thresholds and may follow a different empirical distribution, using the same bin boundaries for both can hide relevant structure. For these reasons, we aim to determine the activation and cancellation bins in a data-driven way, such that the resulting bins

are informative about the cost category of a scenario. (We assess the quality of the rules learned from these bins, not their application to new data; see the closing remark of Section 4.) The total cost binning could in principle be improved in a similar way, however this work focuses on the threshold side for now. The threshold values are the direct input for the countermeasure rules and are expected to be more influential on the interpretability and quality of the resulting rule system than the cost binning. The following subsections introduce two approaches for binning. We first consider an evolutionary binning method based on a genetic algorithm. This provides a flexible heuristic for searching over possible bin boundaries. Afterwards, we show that the ordered structure of the threshold values allows the problem to be solved exactly by dynamic programming. Finally, we compare and discuss these binning results.

5.1 Genetic algorithm threshold binning

As a first approach, we use a genetic algorithm to improve the binning of threshold values. Genetic algorithms are heuristic optimisation methods inspired by biological evolution [22, 24]. Instead of constructing a solution directly, they start with a population of candidate solutions and improve this population over several generations. In each generation, candidates with higher fitness are more likely to survive and produce new candidates. Through selection, crossover, and mutation, the algorithm explores the search space and gradually moves towards better solutions. In our setting, the candidates are possible binnings, and the idea is to let the algorithm "evolve" bin boundaries that are more informative than equally spaced bins. For a fixed number of bins K, one candidate binning is represented by $K-1$ cut-points

$$0 < c_1 < c_2 < \cdots < c_{K-1} < 0.3,$$

which divide the threshold interval $(0, 0.3)$. The upper bound 0.3 comes from the recalibrated threshold sampling range introduced in Section 4.1. Once a candidate set of cut-points is chosen, all threshold values of a simulation are assigned to the resulting bins. The quality of this candidate binning is then measured by a fitness function, which will be introduced shortly. For a fixed number of bins K, the genetic algorithm proceeds as follows:

1. **Initialisation:** Generate an initial population of 50 random candidate binnings. Each candidate consists of $K-1$ randomly chosen cut-points, which are sorted to satisfy the required ordering.

2. **Fitness evaluation:** Evaluate each candidate using a fitness function.

3. **Elitism:** Copy the two candidates with the highest fitness into the next generation. This ensures that the best solutions found so far are not lost and makes the sequence of best solutions found by the GA monotonically increasing.

4. **Selection:** Select parent candidates using tournament selection. In each tournament, three candidates are drawn at random and the one with the highest fitness is selected as a parent.

5. **Crossover:** Generate a child candidate by blending the cut-points of two parents. For each cut-point, the child receives a weighted average

   $$c = wp_1 + (1-w)p_2,$$

   where p_1 and p_2 are the corresponding parent cut-points and $w \sim U(0,1)$.

6. **Mutation and repair:** With probability 0.5, Gaussian noise with standard deviation 0.03 is added to each cut-point. Since this may violate the ordering or the allowed interval, the candidate is repaired by keeping the values within $(0, 0.3)$, sorting the cut-points, and enforcing a minimum distance of 0.001 between neighbouring cut-points.

7. **Iteration:** Repeat selection, crossover, mutation, and repair until 48 new candidates have been generated. Together with the two best candidates, they form the next generation. The procedure is repeated for 100 generations, and the best candidate found during the search is returned.

In short, we use a real-coded, generational genetic algorithm with elitism, tournament selection, blend crossover, and Gaussian mutation. It remains to provide an assessment function for determining the fitness of a candidate binning.

5.1.1 Fitness function

The binning can already be used to recognise some structural information in the simulation data. For this purpose, the bins should be chosen so that they preserve information about the costs of a scenario. In particular, knowing that the bin of a threshold value helps to distinguish whether the corresponding scenario leads to low, medium, or high total cost makes the binning especially useful. If the bins do not help distinguish between these cost categories, then the discretisation has added little beyond making the data technically usable for rule mining. Formally, we use mutual information as the fitness function. Mutual information measures how much knowing one variable reduces the uncertainty about another variable. Following the standard definitions from information theory [11, 36], entropy and mutual information are defined as follows.

Definition 5.1 (Entropy)**.** *For a discrete random variable X with probability distribution $p(x)$, the entropy is defined as*

$$H(X) = -\sum_{x} p(x) \log p(x). \tag{10}$$

Entropy measures the uncertainty of a random variable. If X always takes the same value, then there is no uncertainty and $H(X) = 0$. If X is spread more evenly across several possible values, the entropy (10) is larger.

Definition 5.2 (Conditional entropy)**.** *For two discrete random variables X and Y, the conditional entropy of X given Y is defined as*

$$H(X \mid Y) = -\sum_{y} p(y) \sum_{x} p(x \mid y) \log p(x \mid y). \tag{11}$$

Conditional entropy measures the remaining uncertainty about X after observing Y. In our setting, we use C for the cost category of a simulation,

$$C \in \{\text{low}, \text{medium}, \text{high}\},$$

and B for the bin label assigned to a threshold value. Before looking at any threshold information, the uncertainty about the cost outcome is described by $H(C)$. After applying a candidate binning, each threshold value is assigned to a bin. If knowing B helps predict C, then the binning contains useful information for rule mining. The remaining uncertainty about C after observing B is therefore according to (11)

$$H(C \mid B) = -\sum_{b} p(b) \sum_{c} p(c \mid b) \log p(c \mid b).$$

If the cost distribution inside each bin is almost the same as the overall cost distribution, then knowing the bin does not help much. In this case, $H(C \mid B)$ is close to $H(C)$. If, however, some bins mostly contain low-cost scenarios and other bins mostly contain high-cost scenarios, then knowing the bin strongly reduces the uncertainty about the cost category.

Definition 5.3 (Mutual information)**.** *For two general random variables X and Y, mutual information is defined as*

$$I(X;Y) = H(X) - H(X \mid Y).$$

Equivalently, mutual information can be written as

$$I(X;Y) = \sum_x \sum_y p(x,y) \log \left(\frac{p(x,y)}{p(x)p(y)} \right).$$

This form compares the observed joint distribution $p(x, y)$ with the joint distribution that would be expected if X and Y were independent. If the two variables are independent, then $p(x, y) = p(x)p(y)$ for all x and y, and the mutual information is zero. Applied to our variables, the mutual information between the cost category C and the bin label B is

$$I(C;B) = H(C) - H(C \mid B).$$

It measures how much uncertainty about the cost category is removed by knowing the bin label. Hence, a larger value of $I(C; B)$ means that the binning is more informative about the cost outcome. For the genetic algorithm, this gives a natural fitness criterion. A candidate binning is considered good if the resulting bin labels carry a large amount of information about the cost category. The fitness of a candidate binning is therefore given by

$$\text{fitness} = I(C;B).$$

5.1.2 Optimisation problem

We can now formulate the binning task as an optimisation problem. For a fixed number of bins K, let $\mathcal{P}_K$ denote the set of all admissible binnings of the threshold interval $(0, 0.3)$ into K bins. Each $\pi \in \mathcal{P}_K$ corresponds to a set of ordered cut-points

$$0 < c_1 < c_2 < \cdots < c_{K-1} < 0.3.$$

Applying such a binning to the simulated threshold values gives a discrete bin label B_π for each scenario. The cost category of the same scenario is denoted by C. For a fixed number of bins, the goal is to find the binning that is maximally informative about the cost category. We therefore solve

$$\pi_K^* = \arg\max_{\pi \in \mathcal{P}_K} I(C; B_\pi).$$

Using the identity

$$I(C;B_\pi) = H(C) - H(C \mid B_\pi),$$

this is equivalent to

$$\pi_K^* = \arg\min_{\pi \in \mathcal{P}_K} H(C \mid B_\pi),$$

since $H(C)$ is fixed for a given dataset and does not depend on the binning. In other words, we search for bin boundaries that make the cost distribution inside the resulting bins as informative as possible. The genetic algorithm described above approximates this optimum heuristically.

5.1.3 Worked example for CM1 and $\beta = 0.35$

We illustrate the genetic algorithm threshold binning on the first countermeasure and $\beta = 0.35$, using $K = 5$ bins. Table 5 shows the previous equal-width baseline, in which the activation and cancellation thresholds are pooled and binned by the same fixed cuts. Table 6 replaces these fixed cuts with evolved bin boundaries, still pooling activation and cancellation. Tables 7 and 8 instead evolve separate bin boundaries for the activation and the cancellation threshold.

Table 5: CM1 equal-width, activation and cancellation pooled ($\beta = 0.35$). 5 equal bins over $(0, 0.3)$ (width 0.06); cuts $(0.0600, 0.1200, 0.1800, 0.2400)$. $I(C; B) = 0.140$. $n = 14708$, i.e. 2 times 7354 selected simulations.

	low	med	high	row total
bin 0	2523	1885	1856	6264
bin 1	2886	2323	643	5852
bin 2	181	1977	3	2161
bin 3	0	412	0	412
bin 4	0	19	0	19
col. total	5590	6616	2502	14708

Table 6: CM1 GA binning, activation and cancellation pooled ($\beta = 0.35$). 4 evolved cut-points $(0.0749, 0.1122, 0.1308, 0.1520)$. $I(C; B) = 0.157$. $n = 14708$ i.e. 2 times 7354 selected simulations.

	low	med	high	row total
bin 0	3368	2518	2232	8118
bin 1	1857	1410	262	3529
bin 2	329	735	8	1072
bin 3	36	835	0	871
bin 4	0	1118	0	1118
col. total	5590	6616	2502	14708

Table 7: CM1 activation, GA binning ($\beta = 0.35$). 4 evolved cut-points $(0.0824, 0.1130, 0.1345, 0.2260)$. $I(C; B) = 0.295$. $n = 7354$.

	low	med	high	row total
bin 0	1424	523	1106	3053
bin 1	1099	609	140	1848
bin 2	258	654	5	917
bin 3	14	1478	0	1492
bin 4	0	44	0	44
col. total	2795	3308	1251	7354

Table 8: CM1 cancellation, GA binning ($\beta = 0.35$). 4 evolved cut-points $(0.0428, 0.0722, 0.1302, 0.2000)$. $I(C;B) = 0.087$. $n = 7354$.

	low	med	high	row total
bin 0	1284	1139	962	3385
bin 1	905	916	273	2094
bin 2	601	961	16	1578
bin 3	5	277	0	282
bin 4	0	15	0	15
col. total	2795	3308	1251	7354

Replacing equal-width bins with evolved cuts while still pooling activation and cancellation already raises the mutual information from $I(C;B) = 0.140$ to 0.157. This increase is moderate, but it shows that the choice of binning method can change how cost-relevant structure is extracted from the data. As shown later for other transmission settings, the effect of the binning method alone can be much larger.

Further improvements can be reached in this example by separating the two threshold types: binning the activation threshold on its own reaches $I(C;B) = 0.295$, roughly twice the pooled value, whereas the cancellation threshold on its own reaches only $I(C;B) = 0.087$. This suggests that, for $\beta = 0.35$, the timing of activating CM1 carries substantially more information about the resulting cost category than the timing of cancelling it again. Pooling both quantities into a single binning, as done previously, averages away this difference and understates how informative the activation threshold actually is.

A second observation concerns the last two bins in each table. In every case, the last two bins almost exclusively contain medium cost; specifically, these are simulations in which no countermeasure was activated. It is not surprising that equal-width binning creates two bins that carry essentially the same information and could therefore be combined. However, the same pattern can also be recognised for the GA binning: for both activation and cancellation, the last two bins almost carry the same information and could therefore have been merged in the first place.

This happens because a pure bin like bin 4 contributes positively to the mutual information, since it almost perfectly predicts the cost category, even though it does so based on very little supporting data. Splitting off a small, pure bin like this can never decrease the mutual information, so the genetic algorithm has no incentive to avoid it, and $I(C;B)$ will keep increasing, or at best plateau, as more bins are allowed. This shows that maximising mutual information alone cannot tell us how many bins to use, since it does not penalise binnings that only add sparsely populated bins. We therefore need a criterion that weighs the informativeness of a binning against its complexity when choosing the number of bins.

5.1.4 Akaike information criterion

The Akaike Information Criterion (AIC) is a model-selection criterion [2, 9] that trades off goodness of fit against model complexity.

Definition 5.4 (Akaike Information Criterion)**.** *For a model with maximised likelihood L and k free parameters, the Akaike Information Criterion is defined as*

$$\mathrm{AIC} = 2k - 2\ln L. \tag{12}$$

Lower values are preferred. In our setting, the "model selection" is the choice of the number of bins, K.

Likelihood model for a fixed binning. Consider a candidate binning π with K bins. Each simulation gives one (bin, cost) pair (b_i, j_i), where $b_i \in \{1, \dots, K\}$ is the assigned bin label and $j_i \in \{1, \dots, J\}$ is the observed cost class, with $J = 3$ in our case. Once the binning is fixed, we treat the bin labels as given and assume that the cost classes within each bin follow a multinomial model. Equivalently, each bin has its own categorical distribution over the cost classes:

$$\mathbb{P}(C = j \mid B = b) = p_{bj}, \qquad \sum_{j=1}^{J} p_{bj} = 1.$$

Thus each bin has its own distribution over low, medium, and high cost. If n_{bj} is the number of observations in bin b with cost class j, and $n_b = \sum_j n_{bj}$ is the total number of observations in bin b, the conditional likelihood is

$$L(\pi; p) = \prod_{b=1}^{K} \prod_{j=1}^{J} p_{bj}^{n_{bj}}. \tag{13}$$

Maximising this likelihood (13) under the constraints $\sum_j p_{bj} = 1$ gives the empirical frequencies

$$\hat{p}_{bj} = \hat{p}(j \mid b) = \frac{n_{bj}}{n_b}.$$

This means that the MLE for the probability of a cost class inside a bin is simply the observed share of that cost class in the bin.

Connection to entropy. Substituting the MLE into the likelihood gives the maximised log-likelihood

$$\ln L(\pi) = \sum_{b=1}^{K} \sum_{j=1}^{J} n_{bj} \ln \hat{p}(j \mid b),$$

with the usual convention that terms with $n_{bj} = 0$ contribute zero. Using the empirical joint distribution

$$\hat{p}(b, j) = \frac{n_{bj}}{n},$$

this can be written as

$$\ln L(\pi) = n \sum_{b=1}^{K} \sum_{j=1}^{J} \hat{p}(b, j) \ln \hat{p}(j \mid b).$$

Interestingly, this sum is exactly the negative empirical conditional entropy. Therefore

$$\ln L(\pi) = -n \, \hat{H}(C \mid B).$$

AIC in terms of mutual information. For this conditional multinomial model, each of the K bins has a categorical distribution over J cost classes. Since each such distribution has $J - 1$ free probabilities, the number of free parameters is

$$k = K(J - 1).$$

Substituting the entropy form of the log-likelihood into AIC (12) gives

$$\mathrm{AIC}(\pi) = 2K(J-1) + 2n\,\hat{H}(C \mid B).$$

Using

$$\hat{H}(C \mid B) = \hat{H}(C) - \hat{I}(C;B),$$

we obtain

$$\mathrm{AIC}(\pi) = \underbrace{2K(J-1)}_{\text{complexity}} + \underbrace{2n\,\hat{H}(C)}_{\text{constant}} - \underbrace{2n\,\hat{I}(C;B)}_{\text{fit}}. \tag{14}$$

This form (14) explains the composition of the criterion in our case. The middle term depends only on the simulation data set and therefore is a constant. The last term rewards the same mutual information that we use as the GA fitness. The first term is the complexity penalty: adding one more bin adds $J-1$ free probabilities, so for $J=3$ each additional bin increases the AIC penalty by 4.

Selecting the number of bins. The procedure is therefore two-stage. First, for each allowed number of bins $K \in \{2,\dots,N\}$, we run the genetic algorithm with K fixed and keep the best binning found,

$$\pi_K^\star = \arg\max_{\pi \in \mathcal{P}_K} \hat{I}(C;B_\pi).$$

Second, we compare these candidates by AIC and choose

$$K^\star = \arg\min_{K \in \{2,\dots,N\}} \mathrm{AIC}(\pi_K^\star).$$

Since $2n\,\hat{H}(C)$ is constant across the candidates, this is equivalent to minimising

$$2K(J-1) - 2n\,\hat{I}(C;B_{\pi_K^\star}).$$

Thus AIC does not replace mutual information as the fitness for choosing cut points. It adds the missing step: among the best binnings found for different values of K, it selects the one where the gain in information still justifies the additional binning complexity.

5.1.5 Applying AIC

We computed the AIC for the example from Section 5.1.3 across $n_{\text{cuts}} \in \{1,\dots,9\}$ on CM1's activation threshold. The results are shown in Table 9; the distributions for every n_{cuts} are given in the Supplementary Material Volume [20], section A.1.3.

Looking at the distribution of the AIC-optimal binning ($n_{\text{cuts}} = 9$), shown in Table 10, reveals two separate problems. Even with the AIC penalty in place, the selected binning still contains redundant bins: bins 7–9 are almost pure *medium*-cost bins and could be collapsed into one. Doing exactly that, dropping the last two cuts and treating the remaining seven cuts as an 8-bin solution, gives $I(C;B) = 0.31290$ and AIC $= 10555.02$. This is not just better than the GA's own 8-cut solution ($I(C;B) = 0.31183$, AIC $= 10574.8$); it is better than *every* binning the GA found at any n_{cuts} in Table 9, including the nominal best $n_{\text{cuts}} = 9$ solution. This indicates the usage of an inappropriate version of a GA: We were manually able to find a solution, which is better than the solution provided by the implemented genetic algorithm. Indeed, for $n_{\text{cuts}} = 9$ the fitness in the population last improved at generation 37 and then remained flat for the final 63 generations. In other words, the search stopped improving long before the generation budget was exhausted.

Table 9: GA output across $n_{\text{cuts}} \in \{1, \ldots, 9\}$ for CM1's active threshold ($\beta = 0.35$). AIC selects $n_{\text{cuts}} = 9$ (10 bins).

source	$I(C;B)$	AIC	cut-points
1	0.2247	11827	(0.120)
2	0.2779	11049	(0.087, 0.128)
3	0.2948	10804	(0.082, 0.113, 0.134)
4	0.2949	10807	(0.082, 0.113, 0.134, 0.226)
5	0.3039	10678	(0.063, 0.088, 0.113, 0.134, 0.222)
6	0.3088	10611	(0.055, 0.088, 0.113, 0.128, 0.144, 0.215)
7	0.3117	10571	(0.055, 0.081, 0.098, 0.114, 0.128, 0.144, 0.280)
8	0.3118	10574	(0.055, 0.082, 0.098, 0.113, 0.128, 0.1444, 0.227, 0.273)
9	**0.3133**	**10556**	(0.055, 0.081, 0.092, 0.108, 0.120, 0.130, 0.145, 0.214, 0.258)

Table 10: CM1 active ($\beta = 0.35$), 9 evolved cut-points $(0.0554, 0.0810, 0.0920, 0.1080, 0.1204, 0.1303, 0.1458, 0.2141, 0.2580)$; $I(C;B) = 0.313$ vs. uniform baseline 0.297 (+5%); AIC= 10556.7. Actual counts $n = 7354$.

	low	med	high	row total
bin 0	518	167	606	1291
bin 1	871	344	489	1704
bin 2	387	193	99	679
bin 3	574	294	45	913
bin 4	317	324	10	651
bin 5	97	325	2	424
bin 6	31	496	0	527
bin 7	0	1083	0	1083
bin 8	0	77	0	77
bin 9	0	5	0	5
col. total	2795	3308	1251	7354

The second problem is that even though bins 7–9 can be merged and therefore should not be rewarded, $n_{\text{cuts}} = 9$ is still the AIC-best solution in Table 9. This suggests that AIC's complexity penalty is too weak to rule out such bins. A stricter penalty is therefore required. This may also be an artefact of an early convergence, though.

5.1.6 Outlook: fixing the local pptima and the Weak AIC penalty

The analysis above showed two separate problems. The first problem is that the stochastic nature of a genetic algorithm makes it impossible to *guarantee* the determination of the optimal binning for a given n_{cuts}; a hand-edited version of its own solution beat it outright. One possible response would be to improve the GA itself. In the implemented version of the genetic algorithm, mutation adds Gaussian noise with standard deviation 0.03 to the cut-points. Since the entire threshold interval is only $(0, 0.3)$, this is a rather coarse mutation step and may cause the algorithm to jump over small but relevant improvements. Standard remedies all aim at better exploration: running the GA from many random seeds and keeping the best, increasing the

population size to maintain diversity for longer, reducing or adapting the mutation scale once the population stagnates, injecting diversity through periodic restarts, or combining the GA with a local-search step after each generation in a memetic approach. These changes may help in practice, but they still offer no guarantee. They may simply miss the optimum less often. The general approach of the genetic algorithm will be still a stochastic one.

A more direct response is therefore to replace the heuristic search by an exact algorithm. Because the binnings considered here are *contiguous*, meaning that each bin is an interval of the sorted threshold values, and because the objective decomposes additively across bins, the optimal binning for a fixed n_{cuts} can be found exactly in polynomial time by dynamic programming (section 5.2).

Second, even a perfectly optimised binning would still be scored by AIC, and in the standard setting AIC's penalty is too weak to rule out the redundant bins found above. A stricter complexity penalty is therefore needed. For this reason, the Bayesian Information Criterion (BIC) is introduced alongside dynamic programming in the following section. Alternatively, a stronger weight within the AIC setting might yield similar results.

5.2 Dynamic programming binning

Section 5.1.6 motivated dynamic programming (DP) as an alternate binning approach: rather than searching for good cut-points heuristically, as the genetic algorithm does, DP solves the binning problem exactly. Dynamic programming is a general technique for optimisation problems that can be broken into smaller subproblems [7,10], where the solution to a larger instance is assembled from the (already solved) solutions to smaller instances. In this way, we can compute for a fixed number of bins K the cut-points maximising $I(C;B)$ deterministically.

5.2.1 Applicability of dynamic programming

From Section 5.1.2, maximising $I(C;B)$ for a fixed number of bins K is equivalent to minimising the conditional entropy $H(C \mid B)$, since $I(C;B) = H(C) - H(C \mid B)$ and $H(C)$ does not depend on the binning. This gives the conditional entropy as

$$H(C \mid B) = \sum_{b=1}^{K} \frac{n_b}{n} H(C \mid B = b) = \frac{1}{n} \sum_{b=1}^{K} \text{score}(b).$$

Here, $\text{score}(b) = -\sum_{j=1}^{J} n_{bj} \ln(n_{bj}/n_b)$ depends only on the observations that fall into bin b. Two structural properties make DP applicable here.

1. **Decomposability.** Once we know which observations are inside a bin, we can compute that bin's contribution to $H(C \mid B)$ without knowing where the other bins are. The score of one bin does not change when we move cut-points elsewhere, as long as the observations inside that bin stay the same. This makes the problem decomposable.

2. **Contiguity.** The threshold values are ordered, and bins are intervals covering this ordered sequence. After sorting the observations by their threshold value, every admissible binning is therefore just a partition of this sorted sequence into consecutive blocks. A bin cannot contain two observations while skipping another observation that lies between them in threshold value.

DP searches over the entire space $\pi \in \mathcal{P}_K$ of possible binnings. For K bins and N possible split positions, this gives a polynomial-time algorithm of order $O(KN^2)$.

5.2.2 The algorithm

Setup. We first sort all observations by their threshold value. Since many simulations share the same rounded threshold value, we group equal threshold values together[1]. This gives an ordered sequence of groups $g_0, g_1, \ldots, g_{N-1}$. Choosing K bins means choosing $K-1$ split positions between these groups. Each bin is therefore a consecutive sequence of groups. For any possible block $g_i, \ldots, g_{t-1}$, we compute the entropy contribution that would result from treating this block as one bin:

$$\mathrm{seg}(i,t) = -\sum_{j=1}^{J} n_{i:t,j} \ln\left(\frac{n_{i:t,j}}{n_{i:t}}\right),$$

where $n_{i:t,j}$ is the number of observations with cost class j in the groups $g_i, \ldots, g_{t-1}$, and where $n_{i:t}$ is the total number of observations in this block. Importantly, this value only depends on the observations inside the candidate bin.

DP table. The dynamic program stores the best value for smaller versions of the same problem. We define $\mathrm{dp}[k][t]$ as the minimum total entropy contribution when the first t groups are split into k bins. Thus $\mathrm{dp}[K][N]$ is the value of the optimal K-bin partition of the full sequence. The initial values are

$$\mathrm{dp}[0][0] = 0 \qquad \text{and} \qquad \mathrm{dp}[0][t] = \infty \quad \text{for } t > 0,$$

because zero groups can be split into zero bins, but a positive number of groups cannot be represented using zero bins.

Filling the table. To compute $\mathrm{dp}[k][t]$, we look at the start value of the last bin. Suppose it starts at position i. Then the contribution of the last bin $g_i, \ldots, g_{t-1}$ is $\mathrm{seg}(i,t)$. If everything before i is already optimally split into $k-1$ bins, then it holds

$$\mathrm{dp}[k][t] = \min_{i \in \{k-1,\ldots,t-1\}} \left[\mathrm{dp}[k-1][i] + \mathrm{seg}(i,t)\right].$$

So the result can be determined by trying all possible start values of the last bin and keeping the best one.

Recovering the cut-points. The value $\mathrm{dp}[k][t]$ gives us the optimal score, but not the actual split positions. Therefore, whenever we compute a minimum, we also store which split position achieved it:

$$\mathrm{back}[k][t] = \arg\min_{i} \left[\mathrm{dp}[k-1][i] + \mathrm{seg}(i,t)\right].$$

Starting from $\mathrm{back}[K][N]$, we can then walk backwards through the stored choices. Each stored position gives the start of the last bin in the current subproblem, and therefore one cut-point.

[1] In fact, while there are 7354 selected simulations, meaning 7354 activation thresholds for CM1, there are only $N = 225$ unique threshold values, reducing the search space substantially.

Small example. To make the DP algorithm more intuitive, consider a simplified ordered sequence of six threshold groups. The table below gives for each group $g_0, \ldots, g_5$ the number of low-, medium-, and high-cost observations in that group:

	low	medium	high
g_0	4	0	0
g_1	3	1	0
g_2	0	4	0
g_3	0	3	1
g_4	0	1	3
g_5	0	0	4

Suppose we want to use $K = 3$ bins. The algorithm starts with dp[0][0] = 0. Then it computes the best way to represent prefixes with one bin. Since there is only one bin, the choice is trivial: all groups in the prefix must be placed into the same bin. For example,

$$\text{dp}[1][2] = \text{seg}(0, 2),$$

which corresponds to the single bin $\{g_0, g_1\}$. Next, the algorithm computes the case with two bins. For example, to compute dp[2][4], it considers all possible starts of the second bin:

i	first bin	second bin	total score
1	$\{g_0\}$	$\{g_1, g_2, g_3\}$	9.89
2	$\{g_0, g_1\}$	$\{g_2, g_3\}$	**6.03**
3	$\{g_0, g_1, g_2\}$	$\{g_3\}$	10.40

The best choice is $i = 2$, so

$$\text{dp}[2][4] = 6.03 \qquad \text{and} \qquad \text{back}[2][4] = 2.$$

This means that the best two-bin partition of the first four groups is

$$\{g_0, g_1\} \mid \{g_2, g_3\}.$$

After the two-bin case has been computed, the algorithm moves on to three bins. The final result is dp[3][6], the best way to split all six groups into three bins. Again, DP considers all possible starts of the last bin:

i	best previous 2-bin partition	third bin	total score
2	$\{g_0\} \mid \{g_1\}$	$\{g_2, g_3, g_4, g_5\}$	13.34
3	$\{g_0, g_1\} \mid \{g_2\}$	$\{g_3, g_4, g_5\}$	10.65
4	$\{g_0, g_1\} \mid \{g_2, g_3\}$	$\{g_4, g_5\}$	**9.04**
5	$\{g_0, g_1\} \mid \{g_2, g_3, g_4\}$	$\{g_5\}$	10.65

The best choice is $i = 4$, so

$$\text{dp}[3][6] = 9.04 \qquad \text{and} \qquad \text{back}[3][6] = 4.$$

At this point the process stops, and the optimal score is known. To recover the actual cut-points, we now backtrack. Since back[3][6] = 4, the third bin starts at g_4, so the last bin is $\{g_4, g_5\}$. We then move to the remaining subproblem $(k, t) = (2, 4)$. There it holds

back[2][4] = 2, so the second bin starts at g_2, giving $\{g_2, g_3\}$. The remaining groups form the first bin: $\{g_0, g_1\}$. Thus the recovered optimal partition is

$$\{g_0, g_1\} \mid \{g_2, g_3\} \mid \{g_4, g_5\}.$$

The corresponding cut-points are placed between g_1 and g_2, and between g_3 and g_4. The important point is that the table is filled forward, from small prefixes and few bins to the full problem, while the cut-points are recovered afterwards by walking backward through the stored choices.

5.2.3 Bayesian information criterion

As an alternative to AIC, we also consider the Bayesian Information Criterion (BIC) [9, 35].

Definition 5.5 (Bayesian Information Criterion)**.** *For a model with n observations, maximised likelihood L, and k free parameters, the Bayesian Information Criterion is defined as*

$$\mathrm{BIC} = k \log n - 2 \ln L. \tag{15}$$

For a candidate binning π, we use the corresponding maximised likelihood $L(\pi)$. Using the same argumentation as in 5.1.4 we obtain:

$$\mathrm{BIC}(\pi) = \underbrace{K(J-1)\log n}_{\text{complexity}} + \underbrace{2n\,\hat{H}(C)}_{\text{constant}} - \underbrace{2n\,\hat{I}(C;B)}_{\text{fit}}.$$

The difference to AIC is the complexity penalty. AIC uses $2k$, whereas BIC uses $k \log n$ (see (15). Since $\log n > 2$ for $n > 8$, BIC penalises additional parameters more strongly for the sample sizes considered here[2]. Under AIC, even very small improvements in mutual information can make additional bins look attractive. BIC tends to avoid overly fine discretisations and is therefore better suited to our aim of obtaining interpretable threshold categories.

5.3 Results

The results are organised in three steps. First, the different binning strategies are compared on the most informative example, CM1 at $\beta = 0.2$. This example is used to show why the final procedure is based on dynamic programming rather than the genetic algorithm, and why BIC is preferred over AIC. Second, the selected binnings are summarised across all transmission settings. Finally, the threshold discretisations used for the subsequent rule-mining step are reported in compact form, while the full cut-point lists are deferred to the Supplementary Material Volume [20], section A.1.

5.3.1 Comparison of binning strategies

The comparison includes four binning strategies: equal-width binning (EW), the genetic algorithm (GA), dynamic programming (DP), and dynamic programming with a minimum bin-support floor of 5% (DP+5%). The support floor is applied only to the DP variant and is used as a robustness check against small, high-information bins. For each strategy, the binning selected by AIC and by BIC is reported.

[2]For the CM1 example with $n = 7354$, the BIC penalty per free parameter is $\log(n) \approx 8.90$, compared with 2 for AIC. Thus BIC penalises each additional parameter about 4.45 times as strongly. Since one additional bin adds $J - 1 = 2$ parameters, the penalty per extra bin is 17.80 for BIC instead of 4 for AIC.

The CM1 activation threshold at $\beta = 0.2$ is used as the main comparison case, because it is the threshold with the strongest association to the final cost category in the mild-transmission setting. The results derived based on the different approaches and selected binning methods are shown in Table 11.

Table 11: Selected binnings for CM1 activation at $\beta = 0.2$.

model	criterion	K	$I(C;B)$	AIC	BIC
EW	AIC	37	0.61281	4117.0	4636.6
EW	BIC	37	0.61281	4117.0	4636.6
GA	AIC	28	0.62588	3864.8	4257.9
GA	BIC	14	0.62215	3870.5	4067.1
DP	AIC	22	0.63160	3746.1	4055.0
DP	BIC	10	0.62594	3791.8	3932.2
DP+5%	AIC	9	0.62056	3876.9	4003.3
DP+5%	BIC	8	0.62011	3880.3	3992.7

Two conclusions are visible from this example.

First, BIC selects substantially simpler binnings than AIC. For GA, the selected number of bins drops from $K = 28$ under AIC to $K = 14$ under BIC, while the mutual information decreases only slightly. A similar pattern is obtained for DP, where K drops from 22 to 10. Thus, the stronger BIC penalty removes much of the complexity that AIC accepts for small in-sample gains in mutual information.

Second, a cleaner optimum is obtained with DP than with the GA. This is visible not only from the selected number of bins, but also from the selected cut-points themselves. Under BIC, several GA cuts are still placed in the upper threshold tail, whereas the DP solution stops at 0.1055:

GA+BIC:
(0.0221, 0.0499, 0.0610, 0.0690, 0.0827, 0.0895, 0.1004, 0.1059, 0.1291, 0.1408, 0.2200, 0.2289, 0.2673),

DP+BIC:
(0.0225, 0.0545, 0.0685, 0.0825, 0.0865, 0.0895, 0.0985, 0.1025, 0.1055).

The corresponding distribution tables make the difference more transparent (Table 12). In the GA binning, all observations above the cut at 0.1059 are medium-cost observations, but this pure-medium region is still split into six separate bins. The last three of these bins contain only 23, 34, and 3 observations. In the DP+BIC binning, the same pure-medium tail is kept as one final bin with 3125 observations. Thus, DP achieves slightly higher mutual information (0.62594 instead of 0.62215) while using fewer cuts and avoiding the redundant upper-tail split.

Table 12: Cost-category counts for the BIC-selected CM1 activation binnings at $\beta = 0.2$.

	GA+BIC					DP+BIC			
bin	low	med	high	total	bin	low	med	high	total
1	4	4	4	12	1	4	4	4	12
2	1104	25	198	1327	2	1401	27	226	1654
3	709	21	54	784	3	948	42	44	1034
4	536	23	18	577	4	927	37	0	964
5	927	37	0	964	5	148	126	0	274
6	193	258	0	451	6	45	132	0	177
7	1	698	12	711	7	1	589	1	591
8	0	185	135	320	8	0	203	47	250
9	0	1286	0	1286	9	0	91	99	190
10	0	464	0	464	10	0	3125	0	3125
11	0	1315	0	1315					
12	0	23	0	23					
13	0	34	0	34					
14	0	3	0	3					
total	3474	4376	421	8271	total	3474	4376	421	8271

The 5% support floor further reduces the number of bins in the selected DP+BIC binning from $K = 10$ to $K = 8$, with only a small loss in mutual information (0.62594 to 0.62011). The floored version is therefore used for the actual rule mining: it preserves nearly the same information while enforcing a minimum amount of support for every selected bin, thereby avoiding small bins such as bin 9. This is assumed to be beneficial for rule mining. For the CM1 cancellation threshold at the same β, the mutual information is much smaller across all models, as shown in Table 13. This indicates that, in the mild-transmission setting, the activation timing of CM1 carries much more information about the cost category than the cancellation timing.

Table 13: Selected binnings for CM1 cancellation at $\beta = 0.2$.

model	criterion	K	$I(C;B)$	AIC	BIC
EW	AIC	37	0.17910	11291.4	11810.9
EW	BIC	34	0.17825	11293.5	11770.9
GA	AIC	15	0.17861	11211.6	11422.2
GA	BIC	9	0.17669	11219.2	11345.6
DP	AIC	16	0.18263	11149.1	11373.7
DP	BIC	7	0.17804	11189.0	11287.3
DP+5%	AIC	10	0.16532	11411.3	11551.8
DP+5%	BIC	7	0.16401	11421.0	11519.3

5.3.2 Patterns across transmission settings

After the method comparison, the final selection rule is fixed to DP+5% with BIC. The selected number of bins and the corresponding mutual information values are summarised in Table 14.

The full comparison tables for all models and criteria are provided in the Supplementary Material Volume [20], section A.1.2.

The table 14 shows that the informative threshold variables change with the transmission setting. For $\beta = 0.2$, the CM1 activation threshold is by far the most informative threshold. As β increases, the information carried by CM1 decreases, while CM2 and CM3 become more relevant. CM2 is most informative in the intermediate transmission settings, especially at $\beta = 0.35$, whereas CM3 becomes more informative only in the high-transmission case.

Table 14: Final DP+5%+BIC binning summary.

β	CM	K_{act}	$I_{\text{act}}(C;B)$	K_{cancel}	$I_{\text{cancel}}(C;B)$
0.20	CM1	8	0.62011	7	0.16401
0.20	CM2	8	0.15543	9	0.22802
0.20	CM3	4	0.01755	5	0.05370
0.35	CM1	9	0.31515	8	0.09188
0.35	CM2	9	0.38975	8	0.32280
0.35	CM3	4	0.04018	6	0.11875
0.50	CM1	8	0.15479	6	0.06212
0.50	CM2	9	0.28223	7	0.26678
0.50	CM3	4	0.04159	6	0.08024
0.80	CM1	6	0.06027	3	0.02761
0.80	CM2	7	0.20957	7	0.12263
0.80	CM3	5	0.14200	7	0.08149

Activation thresholds are generally more informative than cancellation thresholds for CM1. For CM2, activation and cancellation are much closer, especially for $\beta = 0.35$ and $\beta = 0.5$. This suggests that, once intermediate countermeasures become relevant, both the timing of introduction and the timing of release contribute to the cost category. CM3 remains weakly informative in the lower-transmission settings, but its information content increases at $\beta = 0.8$.

5.3.3 Final threshold discretisations

According to Table 14, the discretisation used for the rule-mining is selected by DP+5%+BIC. This choice combines three desirable properties. First, DP assures much better than the genetic algorithm approach that the global optimal binning for each fixed number of bins is reached. Second, BIC avoids the overly complex binnings often selected by AIC. Third, the 5% support floor prevents individual bins from being selected based only on very small numbers of observations. The selected numbers of bins and their mutual information values are shown in the main text; the corresponding cut-points and distribution tables are listed in the Supplementary Material Volume [20], section A.1.1.

6 Association rule mining

After constructing the cost labels and discretising the threshold variables, we apply association rule mining to identify interpretable patterns in the simulation data. The goal is not to optimise the thresholds directly, but to find recurring relationships between countermeasure choices and the resulting cost category.

Association rule mining searches for rules of the form

$$X \Rightarrow Y,$$

where X is the condition of the rule and Y is the associated outcome. In our setting, X consists of policy features, such as the activated level of countermeasures or the binned activation and cancellation thresholds. The outcome Y is restricted to exactly one of the three cost labels

$$\{\texttt{Bin=low}\}, \qquad \{\texttt{Bin=med}\}, \qquad \{\texttt{Bin=high}\}.$$

Thus, each retained rule describes which countermeasure patterns are associated with low-, medium-, or high-cost simulation outcomes.

Each selected simulation run is converted into a transaction, i.e. a set of categorical items of the form

$$\texttt{Variable=Category}.$$

For example, a simulation in which two countermeasures are activated may contain items such as

$$\{\texttt{NumCM_Activated=2}, \texttt{T_Active1_Bin=1}, \texttt{T_Cancel1_Bin=0}, \texttt{T_Active2_Bin=3},$$
$$\texttt{T_Cancel2_Bin=2}, \texttt{Bin=low}\}.$$

The association rule mining algorithm is then used to find frequently co-occurring itemsets and convert them into association rules.

We evaluate the rules using support, confidence, and lift, as introduced in [1]. Let $\mathcal{T}$ denote the set of simulated scenarios in the considered rule-mining block. The support of a rule is the share of scenarios in which the condition and the cost outcome occur together:

$$\mathrm{supp}(X \Rightarrow Y) = \mathrm{supp}(X \cup Y) = \frac{|\{t \in \mathcal{T} : X \cup Y \subseteq t\}|}{|\mathcal{T}|}.$$

The confidence is the share of scenarios satisfying X that also have outcome Y:

$$\mathrm{conf}(X \Rightarrow Y) = \frac{\mathrm{supp}(X \cup Y)}{\mathrm{supp}(X)}.$$

Finally, lift compares this confidence with the overall frequency of Y:

$$\mathrm{lift}(X \Rightarrow Y) = \frac{\mathrm{conf}(X \Rightarrow Y)}{\mathrm{supp}(Y)} = \frac{\mathrm{supp}(X \cup Y)}{\mathrm{supp}(X)\,\mathrm{supp}(Y)}.$$

A lift above one means that Y is more likely when X occurs than in the dataset overall, while a lift below one indicates that $\neg Y$ is more likely under condition X.

Rules are retained if their support is at least 0.001 and their confidence is at least 0.3. Lift is not used as an additional filter, but only for ranking and interpretation. After filtering, redundant rules are removed: if a more general rule with the same outcome and at least the same confidence exists, the more specific rule is discarded. This favours shorter and more easily interpretable rules. All measures are reported as descriptive in-sample statistics, as our aim is to assess the quality of the learned rules rather than their predictive accuracy on unseen scenarios.

The rule mining is carried out using a range of different descriptive operators. The block `OnlyNumCM` uses all selected simulations and considers only the number of distinct countermeasure levels activated during each run as the antecedent variable. It therefore gives a broad

overview of how intervention intensity relates to the cost label. The blocks `NumCM1_Activate`, `NumCM2_Activate`, and `NumCM3_Activate` then condition on scenarios with one, two, or three activated countermeasures and use the corresponding activation and cancellation threshold bins. These blocks show which threshold regions are associated with low-, medium-, or high-cost outcomes.

The complete rule tables are reported in the Supplementary Material Volume [20], section A.2. In the main text, we focus on the most interpretable and stable patterns. Rules with high lift but very small counts are treated cautiously, since they may describe rare special cases rather than robust policy patterns.

6.1 Interpretation of extracted policy rules

We now interpret the extracted rules separately for each transmission rate β. For each setting, we first consider the `OnlyNumCM` rules and then use the threshold-specific rules to refine the interpretation. In the rule-mining dataset, `NumCM_Activated` denotes the number of distinct countermeasure level that is activated in a simulation run. The threshold-specific blocks therefore separate scenarios with one, two, or three activated countermeasure levels and analyse the corresponding activation and cancellation threshold bins.

6.1.1 For $\beta = 0.2$

The following table gives an overview of the results of the association rule mining for $\beta = 0.2$ on the total cost and the activated level of countermeasures. Thresholds for support and confidence were set to 0.001 and 0.3, respectively.

Id	rules	support	confidence	lift	count
0	{NumCM_Activated=2} ⇒ {Bin=high}	0.033249	0.689223	13.540532	275
1	{NumCM_Activated=0} ⇒ {Bin=med}	0.377826	1.000000	1.890082	3125
2	{NumCM_Activated=1} ⇒ {Bin=low}	0.417845	0.728039	1.733336	3456

OnlyNumCM. For the mild transmission setting $\beta = 0.2$, the `OnlyNumCM` rules show a clear pattern. Scenarios without activated countermeasures are always assigned to the medium-cost bin, while scenarios with one activated countermeasure are most often low-cost, with confidence 0.728 and 3456 supporting cases. Thus, doing nothing does not lead to the best cost category in this setting, whereas activating the first countermeasure level is frequently associated with favourable outcomes. The rule for two activated countermeasure levels points to high cost and has a large lift of 13.541. However, it is based on only 275 scenarios. It should therefore not be considered as a general argument against two levels of countermeasures, but rather as an indication that fighting a mild-pandemic scenarios with two levels of countermeasures will sometimes lead to high-cost outcomes.

NumCM1 For scenarios with one activated level of countermeasures, the complete threshold-specific rule set is given in the Supplementary Material Volume [20], Table 112. Here activation bins 0–3 of CM1 are strongly associated with low-cost outcomes. In particular, `T_Active1_Bin=0` and `T_Active1_Bin=3` cover many simulations and predict `Bin=low` with confidence close to one, while `T_Active1_Bin=1` and `T_Active1_Bin=2` predict `Bin=low` with confidence one. Higher activation bins no longer show this favourable pattern: `T_Active1_Bin=5` is associated with medium cost, and `T_Active1_Bin=6` is split between medium and high cost.

Thus, high CM1 activation thresholds appear less favourable, especially when combined with low cancellation thresholds.

NumCM2 The corresponding two-level-countermeasure rules are reported in Supplementary Material Volume [20], Table 113. Here, the high-lift low-cost rules are mostly based on very few simulations and should therefore not be seen as robust policy patterns. The more stable signal is the opposite: low CM2 threshold bins are associated with high cost. In particular, `T_Cancel2_Bin=0` predicts `Bin=high` with confidence 0.989 and covers 270 scenarios, while `T_Active2_Bin=0` also points to high cost. Similar high-cost associations occur for low CM1 threshold bins within this block, suggesting that early or low threshold configurations in scenarios reaching two countermeasures are generally unfavourable in the mild transmission setting. The final data contain no scenarios concerning `NumCM3_Activate`.

6.1.2 For $\beta = 0.35$

The following table gives an overview of the results of the association rule mining for $\beta = 0.35$ on the total cost and the activated level of countermeasures. Thresholds for support and confidence were set to 0.001 and 0.3, respectively.

Id	rules	support	confidence	lift	count
0	{NumCM_Activated=3} ⇒ {Bin=high}	0.003264	1.000000	5.878497	24
1	{NumCM_Activated=2} ⇒ {Bin=high}	0.166848	0.473562	2.783835	1227
2	{NumCM_Activated=0} ⇒ {Bin=med}	0.072750	1.000000	2.223096	535
3	{NumCM_Activated=1} ⇒ {Bin=low}	0.377210	0.659848	1.736143	2774
4	{NumCM_Activated=2} ⇒ {Bin=med}	0.182622	0.518333	1.152303	1343
5	{NumCM_Activated=1} ⇒ {Bin=med}	0.194452	0.340152	0.756191	1430

OnlyNumCM. The table shows that for $\beta = 0.35$ one activated countermeasure level is most clearly associated with low cost, with confidence 0.660, lift 1.736, and 2774 supporting scenarios. This suggests that, in many simulations, activating only CM1 provides a favourable balance between epidemic control and intervention burden. The no-countermeasure case is always classified as medium cost, but it covers only 535 scenarios. For two activated levels of countermeasures, the outcome is split between medium and high cost. Medium cost is slightly more frequent in terms of confidence, 0.518, but the high-cost rule has the larger lift, 2.784. Thus, high-cost outcomes are strongly over-represented among scenarios with two levels of countermeasures. The small group of three-level countermeasure scenarios gives an even clearer picture: all 24 scenarios are classified as high cost. Overall, the count-based rules indicate that $\beta = 0.35$ is an intermediate regime: one countermeasure level is often favourable, while activating additional countermeasure levels marks a transition to more costly policy paths.

NumCM1. The threshold-specific rules for one activated countermeasure level are given in the Supplementary Material Volume [20], Table 115. They show a clear pattern in the activation threshold of CM1. Low and moderate activation bins are associated with low cost: `T_Active1_Bin` $= 0, 1, 2, 3$ all predict `Bin=low` with confidence 1.000. The rule for `T_Active1_Bin` $= 4$ still points mostly to low cost, but with lower confidence 0.884. For higher activation bins, the pattern changes. In particular, `T_Active1_Bin` $= 7$ predicts medium cost with confidence 0.933, and `T_Active1_Bin` $= 8$ predicts medium cost with confidence 1.000. This suggests that, when only CM1 is activated, late activation is no longer associated with

the most favourable outcomes. The cancellation threshold of CM1 mainly refines this pattern. Some combinations with intermediate activation bins can still lead to low cost, but the activation threshold is the clearer separator between low- and medium-cost outcomes.

NumCM2. For scenarios with two activated levels of countermeasures, Table 116 shows that the timing of CM2 is especially important. Very early or very persistent use of CM2 is associated with high cost. The clearest example is

$$\{\texttt{T_Cancel2_Bin=0}\} \Rightarrow \{\texttt{Bin=high}\},$$

which has confidence 1.000, lift 2.112, and 536 supporting scenarios. Since a lower cancellation threshold means that CM2 is only released once the control score has fallen further, this points to long or strict CM2 phases as a major source of high total cost. The activation threshold of CM2 supports this finding. $\texttt{T_Active2_Bin} = 0$ is strongly associated with high cost, with confidence 0.988. By contrast, higher CM2 cancellation bins are often associated with medium cost. For example,

$$\{\texttt{T_Cancel2_Bin=4}\} \Rightarrow \{\texttt{Bin=med}\}$$

has confidence 0.959 and covers 541 scenarios. Thus, once CM2 is activated, the release threshold becomes decisive for avoiding high cost. The few low-cost rules have only a small support and should therefore not be considered as policy recommendations.

NumCM3. Only 24 scenarios work with three-level-countermeasures (see Table 117 in the Supplementary Material Volume [20]). All of them are classified as high cost. Consequently, all corresponding threshold-specific rules have confidence 1.000. The lift is 1.000 because there is no variation in the cost label within the subset. This means that the individual threshold rules do not distinguish between different cost outcomes inside the three-countermeasure group. Due to the small number of scenarios in this class, the significance of the conclusions is comparatively low. Overall, the rules for $\beta = 0.35$ suggest that timely activation of the first countermeasure level is usually the most favourable strategy. Activating CM2 is not necessarily harmful, but it moves the system into a more delicate regime where activation and especially cancellation thresholds determine whether the outcome remains medium cost or becomes high cost. Activating all three countermeasures occurs rarely and is consistently associated with high cost in the corrected data.

6.1.3 For $\beta = 0.5$

The following table gives an overview of the results of the association rule mining for $\beta = 0.5$ on the total cost and the activated level of countermeasures. Thresholds for support and confidence were set to 0.001 and 0.3, respectively.

Id	rules	support	confidence	lift	count
0	{NumCM_Activated=3} ⇒ {Bin=high}	0.023168	0.919075	6.123896	159
1	{NumCM_Activated=0} ⇒ {Bin=low}	0.016028	1.000000	2.438877	110
2	{NumCM_Activated=1} ⇒ {Bin=low}	0.276701	0.765726	1.867511	1899
3	{NumCM_Activated=2} ⇒ {Bin=med}	0.353198	0.591220	1.344001	2424

OnlyNumCM. The table above shows that for $\beta = 0.5$ the ordering of intervention intensity is more clear than for $\beta = 0.35$. Scenarios with one activated countermeasure level are mostly associated with low cost:

$$\{\texttt{NumCM_Activated=1}\} \Rightarrow \{\texttt{Bin=low}\},$$

with confidence 0.766, lift 1.868, and 1899 supporting scenarios. This suggests that a single countermeasure level can still provide a favourable balance between epidemic control and intervention burden in this transmission setting. The no-countermeasure case is also associated with low cost, with confidence 1.000, but it covers only 110 scenarios. This rule should therefore be interpreted cautiously: it describes a small group of simulations in which the epidemic trajectory remained mild enough that no countermeasure was triggered. Two activated countermeasure levels are mainly associated with medium cost. The rule

$$\{\texttt{NumCM_Activated=2}\} \Rightarrow \{\texttt{Bin=med}\}$$

has confidence 0.591, lift 1.344, and 2424 supporting scenarios. Thus, activating CM2 often moves the outcome away from the low-cost region, but not necessarily into the high-cost region. In contrast, three activated countermeasure levels are strongly associated with high cost. The rule

$$\{\texttt{NumCM_Activated=3}\} \Rightarrow \{\texttt{Bin=high}\}$$

has confidence 0.919, lift 6.124, and 159 supporting scenarios. Overall, the count-based rules suggest the progression

one CM $\rightarrow$ mostly low cost, two CMs $\rightarrow$ mostly medium cost, three CMs $\rightarrow$ mostly high cost.

NumCM1. The threshold-specific rules for one activated countermeasure level are shown in the Supplementary Material Volume [20], Table 119. As in the previous transmission settings, the activation threshold of CM1 is the clearest separator between low- and medium-cost outcomes. Low and intermediate CM1 activation bins are associated with low cost: $\texttt{T_Active1_Bin} = 0, 1, 2, 3$ all predict `Bin=low` with confidence 1.000. The rule for $\texttt{T_Active1_Bin} = 4$ is also strongly associated with low cost, with confidence 0.976 and 567 supporting scenarios. For higher activation bins, the pattern changes. The rule

$$\{\texttt{T_Active1_Bin=6}\} \Rightarrow \{\texttt{Bin=med}\}$$

has confidence 0.928 and 259 supporting scenarios, while

$$\{\texttt{T_Active1_Bin=7}\} \Rightarrow \{\texttt{Bin=med}\}$$

has confidence 1.000 and 217 supporting scenarios. This indicates that, when only CM1 is activated, late activation is no longer associated with low cost. The cancellation threshold of CM1 is less decisive than the activation threshold. Several cancellation bins still point to low cost, for example $\texttt{T_Cancel1_Bin} = 1$ with confidence 0.843 and $\texttt{T_Cancel1_Bin} = 2$ with confidence 0.837. However, the same cancellation bin can appear in both low- and medium-cost rules, especially for higher threshold values. Thus, the essential recommendation for one countermeasure level is that CM1 should not be activated too late; cancellation choices refine the outcome but do not define the main cost transition.

NumCM2. For scenarios with two activated countermeasure levels, Table 120 in the Supplementary Material Volume [20] shows a more nuanced structure. The `OnlyNumCM` table suggests that two countermeasure levels are mostly associated with medium cost, and the threshold-specific rules explain which CM2 timing choices support this outcome. The most stable medium-cost rules involve higher CM2 activation or cancellation bins. For example,

$$\{\texttt{T_Cancel2_Bin=6}\} \Rightarrow \{\texttt{Bin=med}\}$$

has confidence 0.701 and covers 922 scenarios, while

$$\{\texttt{T_Active2_Bin=5}\} \Rightarrow \{\texttt{Bin=med}\}$$

has confidence 0.749 and covers 860 scenarios. The combined rule

$$\{\texttt{T_Active2_Bin=5}, \texttt{T_Cancel2_Bin=6}\} \Rightarrow \{\texttt{Bin=med}\}$$

is even more specific, with confidence 0.904 and 633 supporting scenarios. This suggests that, once CM2 is activated, later activation and a relatively high cancellation threshold are often compatible with medium cost. At the same time, the table shows clear high-cost patterns. Very low CM2 cancellation thresholds are particularly unfavourable:

$$\{\texttt{T_Cancel2_Bin=0}\} \Rightarrow \{\texttt{Bin=high}\}$$

has confidence 1.000, lift 4.707, and 408 supporting scenarios. Similarly, $\texttt{T_Cancel2_Bin} = 1$ predicts high cost with confidence 0.695 and 239 supporting scenarios. Since lower cancellation thresholds keep CM2 active until the control score has fallen further, these rules indicate that long or persistent CM2 phases are costly. For the activation threshold of CM2 a similar warning can be raised. The rule

$$\{\texttt{T_Active2_Bin=0}\} \Rightarrow \{\texttt{Bin=high}\}$$

has confidence 0.859 and 269 supporting scenarios. Thus, activating CM2 very early is also associated with high total cost. There are also low-cost regions for two countermeasure levels. For example,

$$\{\texttt{T_Active2_Bin=3}, \texttt{T_Cancel2_Bin=6}\} \Rightarrow \{\texttt{Bin=low}\}$$

has confidence 0.950 and 133 supporting scenarios, and

$$\{\texttt{T_Active2_Bin=3}, \texttt{T_Cancel2_Bin=5}\} \Rightarrow \{\texttt{Bin=low}\}$$

has confidence 0.740 and 219 supporting scenarios. These rules show that activating a second countermeasure level is not automatically harmful. Rather, CM2 becomes costly mainly when it is activated too early or cancelled too late. With suitable threshold choices, two countermeasure levels can still remain low or medium cost.

NumCM3. The three-countermeasure block in Table 121 in the Supplementary Material Volume [20] is dominated by high-cost outcomes. The most important high-cost rule is

$$\{\texttt{T_Active3_Bin=0}\} \Rightarrow \{\texttt{Bin=high}\},$$

with confidence 0.924 and 159 supporting scenarios. This is consistent with the `OnlyNumCM` result, where 159 of the three-countermeasure scenarios are classified as high cost. The rule

$$\{\texttt{T_Cancel3_Bin=0}\} \Rightarrow \{\texttt{Bin=high}\}$$

has confidence 1.000 and covers 89 scenarios, giving further evidence for a typically high cost outcome. Many additional high-cost rules have confidence 1.000, for example `T_Active1_Bin` = 1, `T_Active2_Bin` = 0, and `T_Cancel2_Bin` = 2. These rules have varying counts, but they all support the same interpretation: once all three countermeasure levels are activated, the scenario is usually classified as high cost. Some medium-cost rules also appear for three activated countermeasure levels, but they should not be misinterpreted. Their lifts are very large because medium-cost outcomes are rare within the three-countermeasure subset, and most of these rules are supported by only one to five scenarios. Even more prominent medium-cost rules, such as

$$\{\texttt{T_Cancel2_Bin=6}\} \Rightarrow \{\texttt{Bin=med}\},$$

cover only 12 scenarios. Though these threshold combinations may eventually define medium-cost strategies, we can reliably only state that medium-cost exceptions are small compared with the dominant high-cost pattern. Overall, the rules for $\beta = 0.5$ suggest that one activated countermeasure level remains the most favourable intervention strategy, while two countermeasure levels usually lead to medium cost and require careful timing. CM2 should not be activated too early or kept active too long. Activating all three countermeasure levels is strongly associated with high total cost.

6.1.4 For $\beta = 0.8$

The following table gives an overview of the results of the association rule mining for $\beta = 0.8$ on the total cost and the activated level of countermeasures. Thresholds for support and confidence were set to 0.001 and 0.3, respectively.

Id	rules	support	confidence	lift	count
0	{NumCM_Activated=3} ⇒ {Bin=high}	0.097854	0.884226	5.198781	611
1	{NumCM_Activated=0} ⇒ {Bin=low}	0.002402	1.000000	3.447819	15
2	{NumCM_Activated=1} ⇒ {Bin=low}	0.109385	0.785961	2.709851	683
3	{NumCM_Activated=2} ⇒ {Bin=med}	0.497277	0.665025	1.231805	3105

OnlyNumCM. The table above shows that, even in the severe transmission setting $\beta = 0.8$, the activated level of countermeasures follows a clear cost pattern. One activated countermeasure level is associated with low cost, with confidence 0.786, lift 2.710, and 683 supporting scenarios. Thus, low-cost outcomes are still possible, but mainly when the epidemic can be managed with only CM1. The best supported rule in this group is:

$$\{\texttt{NumCM_Activated=2}\} \Rightarrow \{\texttt{Bin=med}\},$$

with confidence 0.665 and 3105 supporting scenarios. This suggests that two countermeasure levels form the typical medium-cost regime for $\beta = 0.8$. By contrast, activating all three countermeasure levels is strongly associated with high cost, with confidence 0.884, lift 5.199, and 611 supporting scenarios. The no-countermeasure rule predicts low cost with confidence 1.000, but covers only 15 scenarios and should therefore be considered as an exceptional case.

NumCM1. If only one countermeasure level is activated (see Table 123 in the Supplementary Material Volume [20]), then low and intermediate CM1 activation thresholds remain favourable. The rules

$$\{\texttt{T_Active1_Bin=0,1,2,3}\} \Rightarrow \{\texttt{Bin=low}\}$$

all have confidence 1.000. The rule for `T_Active1_Bin=4` also points to low cost, with confidence 0.924 and 231 supporting scenarios. The transition occurs at the next activation bin. The rule

$$\{\texttt{T_Active1_Bin=5}\} \Rightarrow \{\texttt{Bin=med}\}$$

has confidence 0.839 and 167 supporting scenarios. This indicates that, under severe transmission, CM1 can still lead to low cost when activated early enough, but delayed activation moves the outcome towards medium cost. The cancellation threshold is less decisive: its rules mostly refine the outcome, while the activation threshold gives the clearer separation.

NumCM2. For two activated countermeasure levels, Table 124 in the Supplementary Material Volume [20] supports the medium-cost pattern from the count-based table. The strongest stable medium-cost rules involve high CM2 threshold bins. In particular,

$$\{\texttt{T_Cancel2_Bin=5}\} \Rightarrow \{\texttt{Bin=med}\}$$

covers 1267 scenarios, and

$$\{\texttt{T_Active2_Bin=5}\} \Rightarrow \{\texttt{Bin=med}\}$$

covers 1230 scenarios. Their combination,

$$\{\texttt{T_Active2_Bin=5}, \texttt{T_Cancel2_Bin=5}\} \Rightarrow \{\texttt{Bin=med}\},$$

has confidence 0.989 and 922 supporting scenarios. This suggests that CM2 can be part of a medium-cost strategy when its activation and release thresholds are placed in suitable higher bins. However, the table also contains high-cost outcomes. Very low CM2 cancellation thresholds are unfavourable: `T_Cancel2_Bin=0` predicts high cost with confidence 0.516 and 129 supporting scenarios, and several more specific rules involving $\texttt{T_Cancel2_Bin} = 0, 1, 2$ reach confidence 1.000. Thus, keeping CM2 active until the score falls very far is costly. Some low-cost regions remain, especially around `T_Active2_Bin=2` and moderate cancellation bins, but the overall role of CM2 at $\beta = 0.8$ is mainly to produce medium-cost rather than low-cost outcomes.

NumCM3. Three countermeasure levels (see Table 125 in the Supplementary Material Volume [20]) usually lead to high-cost outcomes. The prominent rule

$$\{\texttt{T_Active3_Bin=0}\} \Rightarrow \{\texttt{Bin=high}\}$$

covers 409 scenarios, while

$$\{\texttt{T_Cancel1_Bin=0}\} \Rightarrow \{\texttt{Bin=high}\}$$

covers 359 scenarios. Many other high-cost rules have confidence close to or equal to 1.000, but their lift is close to one because high cost is already very common within the situations with three levels of countermeasures. Medium-cost rules also appear, for example rules involving `T_Cancel2_Bin=5` or combinations with `T_Active3_Bin=0`. However, these rules are less prominent than the high-cost pattern and often have much smaller counts. Therefore, activating all three countermeasure levels at $\beta = 0.8$ is a strong indicator of high total cost, while medium-cost cases are exceptions within this block. Overall, the $\beta = 0.8$ rules show that the system shifts towards stronger intervention: one countermeasure level can still be low cost, two countermeasure levels form the main medium-cost regime, and three countermeasure levels mostly indicate high-cost trajectories.

6.2 Summary of rule-mining results

The rule-mining results show a consistent qualitative structure across the transmission settings. A single activated countermeasure level is repeatedly associated with low total cost. This suggests that early or well-placed mild intervention can often reduce infection-related costs without creating excessive economic or compliance burden. Two activated countermeasure levels are more context-dependent: in mild and moderate settings they may indicate difficult or costly trajectories; while for $\beta = 0.5$ and $\beta = 0.8$ they are mainly associated with medium-cost outcomes. Three activated countermeasure levels are consistently associated with high total cost.

The threshold-specific rules refine this picture. For one activated countermeasure level, low- to intermediate CM1 activation bins are generally associated with low cost, while very high activation bins tend to move the outcome toward medium or high cost. This supports the interpretation that delaying the first countermeasure level too much can be costly. For two activated countermeasure levels, the second countermeasure level's activation and cancellation threshold bins become especially important. Depending on the transmission setting, particular CM2 threshold regions separate medium-cost from high-cost outcomes. For three activated countermeasure levels, the rules mainly indicate high-cost trajectories.

Within the selected scenario set, association rule mining identifies interpretable associations between countermeasure thresholds and cost categories. Earlier activation of the first countermeasure is frequently associated with favourable costs, whereas additional countermeasure levels are more frequently associated with medium- or high-cost outcomes. These patterns are conditional on the intervention-release criterion and the chosen cost model; they do not establish that the same associations hold for scenarios with measures still active at the simulation horizon.

7 Discussion, conclusion, and outlook

The main focus of this report was the binning of continuous countermeasure thresholds as a preparatory step for association rule mining. This step is important because the choice of bins directly affects which rules can be found. We have used equal-width binning as a baseline. We have experimented with different data-driven binning methods based on mutual information. The evolution of binning via a genetic algorithm showed that useful bin boundaries are not necessarily equally spaced. However, finding the best solution would have required a large computational effort. Additionally, the question of how many bins should be used was not solvable in a straightforward manner. The dynamic-programming approach is considered as a significant improvement, because it finds the best binning for a fixed number of bins exactly. Together with BIC and the minimum support constraint, it also avoids overly detailed binnings that would be harder to interpret.

This binning approach introduces an important trade-off. In the end, activation and cancellation thresholds are binned separately for each countermeasure level and for each choice of β. This is useful because the parameters under consideration have different empirical distributions and different cost assignments. However, as a consequence the resulting bin labels are not directly comparable across parameters. For example, an activation threshold in bin 3 of CM1 does not represent the same numerical interval as a cancellation threshold in bin 3 of CM2, and the bin widths and bin counts can differ substantially. Therefore, simply comparing bin numbers or subtracting bin labels does not provide a meaningful measure of how far apart two policy thresholds are.

A promising way to address this obstacle is to compute relevant differences before binning, while the variables are still continuous. For each countermeasure, one could first define the activation–cancellation gap

$$G_i = T_i - \widetilde{T}_i,$$

which measures how far the control score must fall after activation before the countermeasure is released again. In addition, differences between activation thresholds, such as $T_2 - T_1$ or $T_3 - T_2$, and the corresponding cancellation-threshold differences can thus supplement the overall picture of countermeasure strategy evaluations. The additionally derived continuous variables could be binned and included in the rule-mining with the same methods as the original model parameters. This would preserve the advantage of data-driven binning while avoiding large errors in the results due to the comparison of values deduced with different underlying discretisations.

Another limitation is that we mainly improved the binning of the threshold variables, but not the binning of the cost values. The binning of activation and cancellation thresholds was optimised with the dynamic-programming approach. The total cost values, however, were simply divided into low, medium, and high cost using a semi-automated heuristic. This made the cost labels more reproducible than a purely manual choice, but they were not optimised in the same way as the threshold bins. Future work may also focus on improving the cost-side binning. One option would be to apply some version of dynamic-programming binning to the continuous total cost values as well. However, other binning methods may also be useful here, or even more effective, since cost binning has different requirements than threshold binning.

In this study, low-cost scenarios are mostly obtained when at most one countermeasure level is activated. This suggests that the calibration of the simulation and cost model is not yet fully balanced. We improved this issue in two ways: first, by recalculating the economic cost component, and second, by shrinking the threshold sampling interval to $(0, 0.3)$, which made activations of CM3 more frequent in the generated data. Nevertheless, it remains difficult to generate scenarios in which activating CM2 or CM3 is clearly worthwhile according to the total cost label. Changing the weights of the cost components can shift this balance, of course.

One way to focus on higher countermeasure levels is to adapt the model to severe epidemic scenarios. Thereafter, late or weak interventions may become substantially more costly. Several model-level improvements could work towards this. In the current total cost function, the health-related cost is mainly based on infections. The compartments H, C, and D are not included as separate cost terms, because in the present model they are strongly correlated with I. Adding them in the same simple way would therefore add little new information. However, these compartments have very different meanings for pandemic management. Hospitalised, critical, and deceased individuals are not just different versions of infection burden. They describe hospital pressure, severe disease, and mortality, which are central aspects of real decision-making. Furthermore, providing care for these patients is usually very resource-intensive. Therefore, future work should include these components in a more meaningful way. One simple idea would be to introduce a maximum hospital-capacity limit. As long as hospital occupancy stays below this limit, the additional hospital costs remain moderate. Once the limit is exceeded, the cost of hospitalisations could increase much more strongly due to the political and psychological effects. This would model the idea that full hospitals create a qualitatively different situation than hospitals that are still below capacity. One could go one step further and also make the critical and deceased compartments depend on hospital overload. If hospitals are full, the number of critical cases and deaths may increase more strongly. Similar extensions could also be considered for C and especially D directly, since mortality is usually one of the most important factors in pandemic management. In the current model, however, the deaths are too closely linked to infections to have much influence on the cost function.

Further upstream, the cost thresholds τ_I and τ_A could be selected more systematically. These thresholds define when infection burden and non-compliance start to contribute to the total cost, so their choice affects all later labels, bins, and rules. Similarly, the control score could be extended beyond the current weighted compartment levels by including derivative information. This would allow the control unit to react not only to the current epidemic burden, but also to expected changes of the situation.

This case study also used only one baseline weighting of the cost components, namely distributing them equally. These weights define how much importance is given to infection burden, non-compliance, and economic cost. Therefore, they can be interpreted as policy preferences. However, this study did not analyse how different weightings would change the resulting rules. Such an analysis would improve our understanding of epidemic management significantly, however. A decision-maker who mainly wants to reduce infections may prefer different threshold regions than a decision-maker who gives more weight to economic cost or non-compliance. Repeating the binning and rule-mining with different cost weightings would show which rules are stable across different preferences and which rules depend on a specific policy perspective.

From the perspective of the authors, the final goal of association rule mining is to provide a solid collection of recommendations for decision-makers. For this, the rules would have to be selected and rewritten in a way that is directly useful in practice. Instead of only reporting many technical rules based on threshold bins, the rule book should focus on a small set of clear recommendations for each transmission setting. It should explain which threshold regions are associated with low, medium, or high cost, and under which conditions stronger countermeasures become necessary or too costly. For practical use, this also means that the rules should be connected more closely to quantities that a decision-maker can observe or estimate in a real situation. In the current model, activation and cancellation depend on the control score, which is useful for the simulation but may be difficult to use directly in real-world decision-making. A future rule book should therefore also translate score-based rules into more accessible interpretations, for example in terms of infection burden, hospital pressure, or non-compliance. Achieving such a rule book, however, first requires the rule-mining step itself to be refined. The rule-mining step still remains relatively high level: the large number of extracted rules suggests that the minimum-support threshold and the later rule-selection procedure still need refinement. Future work should therefore not only mine rules, but also prune, prioritise, and validate them. The given framework can of course also be applied to other infectious diseases. A comparison of the computed rule sets would show how the rule patterns change for diseases with different epidemiological characteristics. However, the resulting rules are not yet a final rule book for decision-makers. They should be seen as an intermediate step towards that goal.

References

[1] R. Agrawal, T. Imieliński, and A. Swami. Mining association rules between sets of items in large databases. In *Proceedings of the 1993 ACM SIGMOD international conference on Management of data*, pages 207–216, 1993.

[2] H. Akaike. A new look at the statistical model identification. *IEEE Transactions on Automatic Control*, 19(6):716–723, 1974.

[3] W. Altaf, M. Shahbaz, and A. Guergachi. Applications of association rule mining in health informatics: a survey. *Artificial Intelligence Review*, 47(3):313–340, 2017.

[4] F. N. G. Andersson and L. Jonung. The covid-19 lesson from sweden: Don't lock down. *Economic Affairs*, 44(1):3–16, 2024.

[5] J. Babić, I. Giray, L. Makaryan, L. Schachinger, C. Kuttler, and J. Draeger. Optimising pandemic response through simulation-based association rule mining. Technical report, Technical University of Munich, 2025.

[6] R. Banerjee, J. Bhattacharya, and P. Majumdar. Exponential-growth prediction bias and compliance with safety measures related to covid-19. *Social Science & Medicine*, 268:113473, 2021.

[7] R. Bellman. *Dynamic Programming*. Princeton University Press, Princeton, NJ, 1957.

[8] R. A. Blair, B. S. Morse, and L. L. Tsai. Public health and public trust: Survey evidence from the ebola virus disease epidemic in liberia. *Social Science & Medicine*, 172:89–97, 2017.

[9] K. P. Burnham and D. R. Anderson. *Model Selection and Multimodel Inference: A Practical Information-Theoretic Approach*. Springer, 2 edition, 2002.

[10] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein. *Introduction to Algorithms*. MIT Press, Cambridge, MA, 4 edition, 2022.

[11] T. M. Cover and J. A. Thomas. *Elements of Information Theory*. Wiley-Interscience, 2 edition, 2006.

[12] L. L. Custode and G. Iacca. Interpretable ai for policy-making in pandemics. In *Proceedings of the Genetic and Evolutionary Computation Conference Companion*, pages 1763–1769, 2022.

[13] Z. Dauzhanov, L. Braun, J. Rei, E. Gül, C. Kuttler, and J. Draeger. Managing an epidemic with costs and non-compliance. Technical report, Technical University of Munich, 2024.

[14] I. Dhaoui, W. Van Bortel, E. Arsevska, C. Hautefeuille, S. T. Alonso, and E. Kleef. Mathematical modelling of covid-19: a systematic review and quality assessment in the early epidemic response phase. *International Journal of Infectious Diseases*, 116:S110, 2022.

[15] J. Draeger. Roadmap to a unified treatment of safety and security. In *Proceedings of the 10th Conference on System Safety and Cyber Security*. IET, 2015.

[16] J. Draeger. Safety, security, and social engineering - thoughts, challenges and a concept of quantitative risk assessment. In J. M. Holder, editor, *A closer look at safety and security*, chapter 3. Nova Science Publishers, Inc., 2020.

[17] J. Draeger and S. Hahndel. Formalized risk assessment for safety and security. *preprint arXiv:1709.00567v2*, 2017.

[18] T. Fischer, T. Gerwald, S. Lajos, S. Woellert, C. Kuttler, and J. Draeger. Modeling the influence of the information domain on countermeasure effectiveness in case of covid-19. In *Journal of Physics: Conference Series*, volume 2514, page 012009. IOP Publishing, 2023.

[19] T. Fischer, T. Gerwald, S. Lajos, S. Woellert, C. Kuttler, and J. Draeger. A Model of the Mutual Influence between Epidemics and Rumors using Covid-19 as an Example. Technical report, Technical University of Munich and IABG mbH, To be published.

[20] N. Granzow, M. Gubitzer, S. Niu, T. Gao, C. Kuttler, and J. Draeger. Supplementary material, effective pandemic management using data-driven threshold binning methods, https://www.researchgate.net/publication/415707184_Supplementary_Material_for_the_Conference_Paper_'Effective_Pandemic_Management_Using_Data-Driven_Threshold_Binning_Methods'. Technical report, Technical University of Munich, Department of Mathematics, 2026.

[21] X. He, F. Min, and W. Zhu. A comparative study of discretization approaches for granular association rule mining. *2013 26th IEEE Canadian Conference on Electrical and Computer Engineering (CCECE)*, pages 1–5, 2012.

[22] J. H. Holland. *Adaptation in Natural and Artificial Systems*. University of Michigan Press, Ann Arbor, 1975.

[23] W. O. Kermack and A. G. McKendrick. A contribution to the mathematical theory of epidemics. *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character*, 115(772):700–721, 1927.

[24] M. Mitchell. *An Introduction to Genetic Algorithms*. MIT Press, Cambridge, MA, 1998.

[25] R. Padmanabhan, N. Meskin, T. Khattab, M. Shraim, and M. Al-Hitmi. Reinforcement learning-

based decision support system for covid-19. *Biomedical Signal Processing and Control*, 68:102676, 2021.

[26] C. Parkinson and W. Wang. Analysis of a reaction-diffusion sir epidemic model with noncompliant behavior. *SIAM Journal on Applied Mathematics*, 83(5):1969–2002, 2023.

[27] T. A. Perkins and G. España. Optimal control of the covid-19 pandemic with non-pharmaceutical interventions. *Bulletin of mathematical biology*, 82(9):118, 2020.

[28] P. A. R. Putri, S. S. Prasetiyowati, and Y. Sibaroni. The performance of the equal-width and equal-frequency discretization methods on data features in classification process. *Sinkron: Jurnal dan Penelitian Teknik Informatika*, 7(4):2082–2098, 2023.

[29] B. Quoc Trung, V. Hoang Minh, N. Thi Hoai Linh, and B. Thi Mai Anh. A novel dynamic programming method for non-parametric data discretization. In *Asian Conference on Intelligent Information and Database Systems*, pages 215–227. Springer, 2024.

[30] S. Ramírez-Gallego, S. García, J. M. Benítez, and F. Herrera. Multivariate discretization based on evolutionary cut points selection for classification. *IEEE transactions on cybernetics*, 46(3):595–608, 2015.

[31] A. Ramponi and M. E. Tessitore. The economic cost of social distancing during a pandemic: an optimal control approach in the svir model: A. ramponi, me tessitore. *Decisions in Economics and Finance*, 48(1):485–516, 2025.

[32] M. Reymond, C. F. Hayes, L. Willem, R. Rădulescu, S. Abrams, D. M. Roijers, E. Howley, P. Mannion, N. Hens, A. Nowé, et al. Exploring the pareto front of multi-objective covid-19 mitigation policies using reinforcement learning. *Expert Systems with Applications*, 249:123686, 2024.

[33] J. Sachs, S. Karim, and L. Aknin et al. The lancet commission on lessons for the future from the COVID-19 pandemic. *The Lancet*, 400:1224–1280, 2022.

[34] A. Salehi and A. Babaei. Network-based analysis of national strategies for covid-19 management. *Iscience*, 28(11), 2025.

[35] G. Schwarz. Estimating the dimension of a model. *The Annals of Statistics*, 6(2):461–464, 1978.

[36] C. E. Shannon. A mathematical theory of communication. *Bell System Technical Journal*, 27(3):379–423, 1948.

[37] S. Sinisterra-Sierra, S. Godoy-Calderón, and M. Pescador-Rojas. Covid-19 data analysis with a multi-objective evolutionary algorithm for causal association rule mining. *Mathematical and Computational Applications*, 28(1):12, 2023.

[38] R. Somyanonthanakul, K. Warin, W. Amasiri, K. Mairiang, C. Mingmalairak, W. Panichkitkosolkul, K. Silanun, T. Theeramunkong, S. Nitikraipot, and S. Suebnukarn. Forecasting covid-19 cases using time series modeling and association rule mining. *BMC medical research methodology*, 22(1):281, 2022.

[39] U. Stańczyk, B. Zielosko, and G. Baron. Discretisation of conditions in decision rules induced for continuous data. *PloS one*, 15(4):e0231788, 2020.

[40] M. Tandan, Y. Acharya, S. Pokharel, and M. Timilsina. Discovering symptom patterns of covid-19 patients using association rule mining. *Computers in biology and medicine*, 131:104249, 2021.

[41] R. N. Thompson. Epidemiological models are important tools for guiding covid-19 interventions. *BMC medicine*, 18(1):152, 2020.

[42] J. Wells, C. Robertson, V. Marmara, A. Yeung, and A. Kleczkowski. Modelling a novel coronavirus (covid-19): A stochastic seir-hcd approach, with real-time parameter estimation and forecasting for scotland. *arXiv preprint arXiv:2012.12390*, 2020.

[43] F. Zurek, L. Schwarzmeier, B. Dionisi Ferrera, C. Kuttler, and J. Draeger. Managing epidemics: Dynamic countermeasures and frictions. Technical report, Technical University of Munich, Department of Mathematics, 2023.