Governing the Unmeasurable: Rules, Discretion, and the Actuarial Convention

October 2026

Ilker Aslan, PhD[1]

Abstract

Actuaries certify numbers that no measurement technology delivers: a one-in-two-hundred annual standard no claims record could support, premium risk factors held constant across the underwriting cycle, and, as machine-learned components enter pricing and reserving, outputs whose basis the certifier cannot reconstruct. This paper treats such numbers as conventions in Keynes's sense and asks how they should be governed. A convention is formalized as a published mapping from a verifiable signal to a binding standard, and rule, discretion and the frozen parameter are compared on time consistency, resistance to capture and auditability. Discretionary revision relaxes the standard where the signal indicates most risk and leaves no record, because the number is the record. Rules trade the frequency of capture for its persistence, so a bound on the mapping's output limits distortion, and the bound's optimal width is where the stop-loss premium of the unbounded adjustment equals the capture-weighted bound. Applied to Solvency II, the framework sorts the symmetric adjustment, the 99.5 per cent standard and the premium risk factors into three cells; EIOPA's daily series since 1991 shows the ±10 corridor binding at three published month-ends in a decade, and widths of 10, 13 and 17 points imply revision-survival probabilities of 0.87, 0.94 and 0.98. A closing section applies the criteria to certifying model outputs that cannot be validated.

Keywords: conventions; fundamental uncertainty; solvency regulation; rules versus discretion; regulatory capture; actuarial standards

JEL codes: G22, G28, D81, E61, K23

1. Introduction

The solvency capital requirement of an insurer established in the European Union corresponds to the value-at-risk of its basic own funds at a confidence level of 99.5 per cent over one year.[2] No

[1]Lecturer, Kalaidos University of Applied Sciences, Zürich. ORCID 0000-0001-9241-9420. Correspondence: ilker.aslan@alumni.unifr.ch. Prepared for the Joint Colloquium of the IAA Sections, Luxembourg, 19–21 May 2027 (AFIR–ERM Section). The paper is informed by a decade of treaty reinsurance underwriting practice; the views expressed are the author's alone. Replication code and data for the numerical and empirical results accompany the paper.

[2]Directive 2009/138/EC, Article 101(3).

actuary defends that quantile as an estimated feature of a portfolio loss distribution. A one-in-two-hundred annual event is, by construction, an event the record of any single undertaking has had almost no opportunity to produce, and the statistical apparatus that would recover it from data does not exist at the portfolio level. The number nevertheless disciplines capital across a continent, coordinates supervisory judgment, and is revised through an identifiable process. It works without being a measurement.

That condition is the ordinary one for standard-setting at the tail, and it is not confined to the headline quantile. The standard formula's premium risk module applies a volatility factor of 14 per cent to general liability business and 17 per cent to non-proportional casualty reinsurance, fixed in an annex and invariant across the underwriting cycle.[3] The symmetric adjustment to the equity capital charge moves the charge by up to ten percentage points according to a published formula whose output nobody claims to be the correct adjustment, only the agreed one.[4] The risk margin rests on a cost-of-capital rate set in legislation. In each case a number that cannot be estimated is set, published, used, and periodically changed.

The profession has an elaborate vocabulary for the technical content of such numbers (calibration, validation, expert judgment) and almost none for their institutional status. The question the calibration vocabulary cannot reach is the one this paper takes up: given that the number is not an estimate, what makes one way of setting and revising it better than another? The question has become pressing for a reason external to the regulatory apparatus. As statistical learning methods enter pricing, reserving, and underwriting, the actuary signing a technical provision or a rate increasingly stands where the supervisor has always stood with respect to the 99.5 per cent standard: certifying a number whose basis cannot be reconstructed from the evidence available to the certifier. The governance problem that regulation solved by long institutional habit becomes one the profession must solve explicitly, at the level of the individual signature.

I take these numbers to be conventions in the sense Keynes (1936, ch. 12; 1937) gave the term: judgments agents coordinate on because coordination is necessary, sustained by practice and by the expectation that others will sustain them, and revised when the practice that supports them gives way. The reading itself is not new. Power (1997) applied it to accounting measurement and Orléan (2014) to financial valuation, and I have argued elsewhere that reinsurance underwriting cycles are driven by the formation and collapse of conventional judgments about loss distributions that the data do not identify (Aslan 2026c). What has not been done, so far as I can find, is to ask the design question that follows. If a standard is a convention, then setting it is an act of institutional design, and designs can be compared.

The comparison runs along three dimensions, each with an established literature behind it. The first is *time consistency*: whether the authority setting the standard would still wish, after the regulated firms have committed capital, to set the standard it announced beforehand. This is the problem

[3] Commission Delegated Regulation (EU) 2015/35, Article 117 and Annex II, as substituted by Commission Delegated Regulation (EU) 2019/981. The segment standard deviations for gross premium risk run from 6.4 per cent (assistance) to 19 per cent (credit and suretyship); Article 117(3) applies an adjustment factor for non-proportional reinsurance of 80 per cent to segments 1, 4 and 5 and 100 per cent elsewhere.

[4] Directive 2009/138/EC, Article 106; Commission Delegated Regulation (EU) 2015/35, Article 172. The adjustment equals one half of the excess of the current equity index level over its equal-weighted daily average across the preceding 36 months, less eight percentage points, and is confined to the interval from -10 to $+10$ percentage points. Directive (EU) 2025/2 widens the interval to ± 13 percentage points and Commission Delegated Regulation (EU) 2026/269 carries the change into Article 172, both with effect from 30 January 2027.

Kydland and Prescott (1977) identified in monetary policy and Barro and Gordon (1983) formalized. The second is *resistance to capture*: how far the standard can be moved by the interests it binds, following Stigler (1971) and the regulatory economics that grew from it. The third is *auditability*: whether an outside party, given what has been published, can reproduce the standard and so tell whether it was set as announced.

The paper's results are six. Under discretionary revision the standard is set as a weighted average of what the signal implies and what the regulated firm has already done, so it is relaxed most where the signal indicates the greatest risk; firms anticipate this and hold less capital in advance (Proposition 1). The mechanism is familiar as forbearance, but its form here is worse: forbearance by non-enforcement leaves a visible gap between the standard and the supervisor's conduct, whereas forbearance by re-parameterization leaves none, because the number is the record of what the standard was. Under capture, a rule reduces the frequency with which the standard can be moved but multiplies the persistence of each movement, and what caps the distortion independently of the revision interval is a bound on the mapping's output (Proposition 2). Auditability and responsiveness are jointly attainable only under a rule, which is why a frozen parameter is at once perfectly auditable and informationally inert (Proposition 3). A rule with an escape clause on a pre-announced threshold that the firm cannot move dominates both pure regimes when the signal omits information the supervisor will later hold, and an escape conditioned on anything the firm controls reproduces discretion in full (Proposition 4). When the signal is a statistic of the regulated population itself, the manipulation wedge rises with that population's Herfindahl concentration (Proposition 5). And the bound has an optimal width of a form actuaries will recognize: the corridor is set where the stop-loss premium of the unbounded adjustment at that retention equals the capture-weighted bound, so that a proposal to widen a corridor is a claim about the signal's tail and about the survival of captured revisions, and can be judged as one (Proposition 6). A calibrated version of the model shows the magnitudes: discretion cuts the capital held in advance by a third, and an escape clause invoked in one period in six lowers the supervisor's expected loss below both pure regimes.

Section 4 applies the framework to three parameters of the European regime that occupy three cells of the resulting taxonomy: the symmetric adjustment, a rule in the full sense; the 99.5 per cent standard, a number with no convention governing it, and rightly so, because it encodes a preference about how much safety to buy and not a proxy for a state of the world; and the premium risk volatility factors, which proxy for a state that is not constant and are revised at occasions no one announces. The record of the regime's revision between 2019 and 2027 behaves as the model says the three cells should, and EIOPA's own daily series of the unbounded adjustment, which runs from 1991, puts Proposition 6 to work: the ± 10 corridor bound at three of the 129 month-ends the regime has published, a ± 13 corridor would have bound at one and a ± 17 corridor at none, and on the long history the three widths are rationalized by revision-survival probabilities of 0.87, 0.94 and 0.98, so that the choice among them turns on a convention about the tail of the kind the paper is about. Section 5 carries the framework to the certification of model outputs the certifier cannot validate and argues that the professional obligation attaching to such a number concerns the governance of the convention and not the accuracy of the number, a claim an actuary can support in place of one no evidence can.

Section 2 fixes the objects and states the scope condition. Section 3 develops the model and the numerical illustration. Section 4 applies it to the European regime and to EIOPA's series; Section

5 to certification. Section 6 takes up objections, including the objection that the exercise amounts to dressing up indexation, and states what would revise the conclusions. Proofs are collected in Appendix A and the calibration in Appendix B.

2. Standards, Conventions, and the Measurement Gap

2.1 What a standard is, and what it is not

Fix terms. A *standard* is a number, published by an institution, that enters a binding requirement: a capital charge, a reserving basis, a rate-filing constraint. A *convention* is the apparatus that produces and revises the standard: the signal it responds to, the mapping from signal to standard, the occasions on which that mapping may itself be changed, and the party who owns it. A standard can exist without a convention, and this is more common than it looks; much of Section 4 turns on the distinction.

The scope condition for everything that follows is that the quantity the standard refers to is not estimable at the precision the standard asserts. I state this as a condition and not as a general claim about actuarial science, because it is false for a great deal of actuarial work and the paper has nothing to say about that part. Attritional motor frequency is estimable. The 99.5th percentile of a one-year change in basic own funds is not, and the reason is not that the data are poor but that the object sits at a quantile the data cannot reach: the record is short relative to the return period, the generating process is not stationary across the record, and the tail is joined to the body by no assumption that is itself testable. Where a quantity is estimable, calibration is the right frame and the machinery of this paper is unnecessary. Where it is not, the number still has to be set, and the question of how it should be set and revised cannot be answered by better estimation.

The distinction between the two situations is the one Keynes (1921) and Knight (1921) drew and Davidson (1991) and Dequech (2011) have sharpened: between a setting in which a probability distribution is available to the agents and one in which it is not, because the generating process is singular, unprecedented, or in transformation. The strongest version of that claim is more than the argument needs. What it needs is that over the relevant horizon no measurement technology delivers the quantity as a point estimate an institution could treat as data. Whether that is because the parameter is poorly identified or because there is no parameter there to identify makes no difference to the design problem; the two readings converge on the same institutional requirement, as I have argued in a related setting (Aslan 2026b, secs. 2.3 and 4.4).

2.2 Conventional judgment

Keynes's account of conventional judgment in Chapter 12 of the *General Theory* is usually read as a description of asset markets, and its application to valuation under uncertainty has been developed in that direction by Orléan (1999, 2014). The feature that matters here is structural rather than psychological. A convention solves a coordination problem in which the thing coordinated upon has no independent anchor: its authority comes from being generally adopted and expected to remain so, not from corresponding to a measured quantity (Lewis 1969). Keynes's own emphasis fell on the fragility that follows, since a convention sustained by nothing but mutual expectation can collapse when the expectation does, and that fragility is the familiar object of the underwriting-cycle literature, in which market-wide judgments about loss distributions harden and then give way

(Cummins and Outreville 1987; Harrington and Danzon 1994; Lamm-Tennant and Weiss 1997; Aslan 2026c). The terms on which reinsurance capacity is offered (attachments, aggregates, wordings) encode the market's confidence in its own models and move on a margin of their own, separately from price (Aslan 2026e).

Regulatory standards are conventions with an institutional sponsor, and the sponsorship changes the analysis in a way that is central to this paper. A market convention has no owner; it moves when enough participants move. A regulatory convention has an identified owner, a published form, and a revision procedure. These features are the levers of institutional design, and the argument of the paper is that they, and not the number, deserve the profession's attention.

The conventions reading does not claim that the numbers are arbitrary; a convention can be well or badly chosen and more or less responsive to the evidence that accumulates. The claim is narrower: its authority does not derive from measurement, and arguments that presuppose it does, in defence of a standard or in attack on it, are misdirected.

2.3 Three regimes

Two properties of a convention organize the field. A convention is *responsive* if the mapping from signal to standard is non-trivial, so that the standard moves when the published signal moves. It is *determinate* if the mapping is single-valued: given the published signal history and the published mapping, the standard that will be set is pinned down.

Cross-classifying gives four cells, three of which are occupied in practice. A *rule* is responsive and determinate: the standard is a published function of a published signal, and anyone can compute it. *Discretion* is responsive and indeterminate: the authority observes the signal, and much else, and sets the standard each period on its own judgment, so that the standard moves with conditions but which standard will be set cannot be derived from what has been published. A *frozen parameter* is determinate and unresponsive: the number is fixed until it is changed by an exercise that is not part of the convention. The fourth cell, unresponsive and indeterminate, would be a standard that neither tracks anything nor can be predicted, and it is not observed, for the sufficient reason that it could not survive contact with the regulated population.

The three regimes are the objects compared in the next section. One preliminary point about what the comparison is not: it is not the question whether a standard should be prudent or lenient. Each regime can be implemented at any level of severity. The question is what happens to the standard over time, under the pressures that act on the institution that sets it.

3. A Model of Convention Governance

3.1 Environment

Three dates, $t \in \{0, 1, 2\}$; a supervisor and a representative regulated firm.

Let $\theta \in \mathbb{R}_+$ denote the capital the firm would need in order to meet the regime's substantive objective. By the scope condition of Section 2.1, θ is neither observed nor estimable. There is a public signal $x \in \mathbb{R}_+$, verifiable by any third party, and a shift $\beta \in \mathbb{R}$ that the supervisor observes at $t = 2$ but that cannot be written into a mapping in advance, with

$$\theta = x + \beta + \varepsilon, \qquad \mathbb{E}[\beta] = 0, \quad \mathrm{Var}(\beta) = \tau^2, \qquad \mathbb{E}[\varepsilon \mid x, \beta] = 0, \quad \mathrm{Var}(\varepsilon \mid x, \beta) = \sigma_\varepsilon^2 > 0, \quad (1)$$

where β and ε are independent of x and of each other. The signal is informative and insufficient; σ_ε^2 does not shrink with the sample, which is the formal content of the claim that the quantity is not estimable. The shift β represents what the supervisor will know at the time the standard is set that no mapping could have anticipated: a liquidity event, a change in the legal environment, a structural break in the relation between the signal and the state. The essential asymmetry is that x is contractible and θ and β are not: a mapping can be written on the signal and cannot be written on the state.

A *standard* $s \in \mathbb{R}_+$ is published at $t = 2$ and requires $k \geq s$, where k is capital the firm chose at $t = 1$ at unit cost $c > 0$. Where $k < s$ the firm must close the gap at $t = 2$, at convex cost $\frac{\psi}{2}(s-k)_+^2$ with $\psi > c$, reflecting that capital raised under a binding constraint is dearer than capital raised in advance.

The supervisor's loss at $t = 2$ is

$$L_S(s; \theta, k, x) \;=\; \tfrac{1}{2}\,(s - \theta)^2 \;+\; \tfrac{\gamma}{2}\,\phi(x)\,(s - k)_+^2, \quad (2)$$

with $\gamma > 0$ and $\phi(\cdot)$ positive and increasing. The first term is prudential error: a standard away from the state in either direction is costly, too low because the firm is under-capitalized and too high because capital is socially scarce. The second term is the cost of a binding shortfall, and ϕ increasing in x encodes the central regularity of the procyclicality literature: forcing an insurer to raise capital or shed risk is most damaging when conditions are worst, because everyone is doing it at once and the assets being shed are the assets everyone is shedding (Danielsson et al. 2001; Plantin, Sapra and Shin 2008; Adrian and Shin 2010). Nothing below turns on the quadratic form beyond tractability; what matters is that the shortfall cost is convex and that its weight rises with the signal.

The timing is that the regime is announced at $t = 0$, the firm chooses k at $t = 1$ knowing the regime but neither x nor β, and at $t = 2$ the signal and the shift realize and the standard is set according to the regime. Under a committed rule the standard is $s = r(x)$ for a published r; the benchmark rule throughout is $r(x) = x$, the conditional expectation of the state given the contractible information, which is the rule a supervisor who could commit and who ignored the shortfall term would choose. Sections 3.2 to 3.4 set $\tau = 0$ to isolate the accommodation mechanism; Section 3.5 restores the shift.

3.2 Discretion relaxes the standard where the signal is worst

Under discretion the supervisor sets s at $t = 2$, observing x and k, with k sunk. With $\tau = 0$ the supervisor's estimate of the state is $\mathbb{E}[\theta \mid x] = x$. The first-order condition for an interior solution with $s > k$ is $(s - x) + \gamma\phi(x)(s - k) = 0$, so that

$$s_D(x, k) \;=\; \frac{x + \gamma\phi(x)\,k}{1 + \gamma\phi(x)}. \quad (3)$$

The discretionary standard is a weighted average of the signal-implied requirement and the firm's

existing capital, with the weight on the latter increasing in the procyclicality term $\gamma\phi(x)$. The standard therefore lies below the signal-implied requirement whenever the firm is under-capitalized relative to the signal,

$$x - s_D(x,k) \;=\; \frac{\gamma\phi(x)}{1+\gamma\phi(x)}\,(x-k) \;>\; 0 \quad \text{for } k < x, \tag{4}$$

and the shortfall is increasing in $\gamma\phi(x)$ and hence in the signal itself. The standard is laxest in the states where the signal says the risk is greatest.

The firm, at $t = 1$, knows (3). Its exposure to a shortfall is $s_D - k = (x-k)_+/(1+\gamma\phi(x))$, so its problem is

$$\min_k \; ck \;+\; \frac{\psi}{2}\,\mathbb{E}\!\left[\frac{(x-k)_+^2}{(1+\gamma\phi(x))^2}\right], \tag{5}$$

with first-order condition $c = \psi\,\mathbb{E}\big[(x-k)_+/(1+\gamma\phi(x))^2\big]$. Under the committed rule $r(x) = x$ the same firm faces the condition $c = \psi\,\mathbb{E}\big[(x-k)_+\big]$.

Proposition 1. *Let $\gamma\phi(x) > 0$ on a set of positive measure. Then the capital chosen under discretion satisfies $k_D < k_R$, and the standard realized under discretion satisfies $s_D(x, k_D) < x$ in every state in which $k_D < x$. The wedge $x - s_D$ is increasing in $\gamma\phi(x)$.*

The proof, in Appendix A, is a comparison of the two first-order conditions; the substance is in (3) and (4). Some remarks on interpretation are in order.

The mechanism does not require a weak or captured supervisor. The supervisor in this model faithfully minimizes a social loss function that correctly includes the cost of forcing deleveraging in a stressed market. Discretion produces the laxer standard because the supervisor does its job with the firm's capital taken as given, and the firm's capital is lower because the firm anticipated exactly that. This is the Kydland–Prescott structure, and its usual implication carries over: a supervisor who weights the shortfall term less heavily reduces the wedge but does not remove it, which is Rogoff's (1985) partial remedy of the conservative central banker transposed to prudential supervision. The same structure has been identified in bank closure policy, where the regulator's ex post reluctance to close weak institutions is anticipated ex ante (Mailath and Mester 1994; Acharya and Yorulmazer 2007), and in the self-interested regulator of Boot and Thakor (1993).

The object being relaxed here is a parameter, not an enforcement decision, and in this respect the setting differs from the forbearance literature it otherwise resembles (Kane 1989). When a supervisor declines to act on a breach, the breach and the inaction are separate facts, and the gap between them can be observed by a legislature, an auditor, or a court. When the supervisor instead re-sets the parameter, there is no gap, because the parameter is the definition of compliance. The firm is in compliance with a standard that moved. Nothing in the published record distinguishes a standard that moved because the state moved from one that moved because the firm did not.

The model's ϕ also gives the signature to look for in practice. Forbearance by re-parameterization is most tempting in the conditions where the prudential case for tightening is strongest, so the diagnostic is not that standards are lax on average but that they are least responsive at the top of the signal's range.

3.3 Capture: rules trade frequency for persistence

Let an interested party (in the applications below, the regulated sector acting through its representative bodies) be able to shift the standard downward by $\delta \geq 0$ at convex influence cost $\frac{a}{2}\delta^2$ incurred on each occasion at which the standard is determined, and let the sector's per-period benefit from a downward shift be $b\delta$ with $b > 0$.

Under discretion the standard is determined every period, so influence is exerted every period on a one-period prize: $\max_\delta \; b\delta - \frac{a}{2}\delta^2$, giving $\delta_D = b/a$ per period. Under a rule revised every T periods, influence is exerted once per T periods on a prize that persists for T periods: $\max_\delta \; Tb\delta - \frac{a}{2}\delta^2$, giving $\delta_R = Tb/a$, which then obtains in each of the T periods.

Taken alone this comparison favours discretion, and presenting the rule as dominant would misstate the result. What rules buy is not immunity but visibility. A rule revision is a single, dated, documented event with a published before and after; a discretionary re-setting is one of a stream of routine determinations. Let the distortion be contested and reversed with probability q_R under a rule revision and q_D under a discretionary determination, with $q_R > q_D$ reflecting that difference in salience. Finally, let the mapping's output be bounded: $\delta \leq \bar{\delta}$.

Proposition 2. *Expected per-period distortion is* $(1 - q_D)\min\{b/a,\ \bar{\delta}\}$ *under discretion and* $(1 - q_R)\min\{Tb/a,\ \bar{\delta}\}$ *under a rule revised every* T *periods. In the absence of a binding bound, the rule yields the lower distortion if and only if*

$$q_R \;>\; 1 - \frac{1 - q_D}{T},$$

so that the rule's advantage is decreasing in the revision interval and vanishes for T *large. Where the bound binds under both regimes, the distortion is* $(1 - q_i)\bar{\delta}$ *and the ranking depends only on salience.*

The result is unwelcome for a naive case in favour of rules, and I think it is the most useful thing in the paper. Durability, the property that makes a rule valuable against time inconsistency, is a liability against capture: it converts a stream of small prizes into one large one and concentrates the interested party's effort on a single occasion where the stakes justify a great deal of effort. The familiar intuition that rules are harder to capture holds only through the salience channel, and the salience channel is weak when revisions are far apart, technical, and conducted in a forum the public does not attend, which describes the revision of actuarial parameters rather exactly. The mechanism, concentration of stakes at infrequent decision points, is a theme of the capture literature from Stigler (1971) through Laffont and Tirole (1991) to the survey of Dal Bó (2006) and the collection of Carpenter and Moss (2014); its application to the durability of a rule appears not to have been drawn.

What survives is the bound. A cap on how far the mapping may move the standard limits the distortion whatever the revision interval and whatever the capture probability, because it does not depend on the political economy at all. This gives a reading of a familiar feature of the European regime that I have not seen stated: the corridor on the symmetric adjustment is usually discussed as a device for limiting the volatility of capital requirements, and it does that, but its more durable function is to cap the prize available to anyone who succeeds in influencing the mapping. The corridor is a capture-control device.

The bound is itself a convention, set and revisable, and the argument therefore recurses: what bounds the bound? The answer is that the recursion terminates in a legislative instrument whose amendment is expensive, so that the location of a parameter in the hierarchy of legal instruments is an element of its governance and not a technicality. Section 4.4 shows the recursion in operation. A further caveat: a bound that binds often is a bound that suppresses information, and the design problem is to set it where it caps the tail of the distribution of adjustments without truncating the body. That problem has a solution of a form actuaries will recognize.

Let the mapping's unbounded output be $u = r(x)$, with a distribution symmetric about zero, and let the published standard be its truncation to the corridor, $s = \max\{-\bar{\delta}, \min\{u, \bar{\delta}\}\}$. Where the bound binds the convention fails to transmit $|u| - \bar{\delta}$ of the signal, at prudential cost $\frac{1}{2}(|u| - \bar{\delta})_+^2$. Capture, by Proposition 2, shifts the standard by $\min\{\delta^*, \bar{\delta}\}$ with probability $1-q$, where $\delta^* = Tb/a$ is the shift the interested party would obtain without a bound, at prudential cost $\frac{1}{2}\min\{\delta^*, \bar{\delta}\}^2$. The supervisor sets $\bar{\delta}$ at $t = 0$ to minimize

$$L(\bar{\delta}) \;=\; \tfrac{1}{2}\,\mathbb{E}\big[(|u| - \bar{\delta})_+^2\big] \;+\; \tfrac{1}{2}\,(1-q)\,\min\{\delta^*, \bar{\delta}\}^2. \tag{6}$$

Proposition 6. *On $[0, \delta^*]$ the loss (6) is strictly convex and its minimizer $\bar{\delta}^*$ is the unique solution of*

$$\mathbb{E}\big[(|u| - \bar{\delta})_+\big] \;=\; (1-q)\,\bar{\delta},$$

that is, the bound is set where the stop-loss premium of the unbounded adjustment at retention $\bar{\delta}$ equals the capture-weighted bound. The optimal bound is increasing in q and widens under a mean-preserving spread of u. The bound is worth having at all if and only if $L(\bar{\delta}^*) < \frac{1}{2}(1-q)\delta^{*2}$, which holds whenever the unconstrained captured shift is large relative to the dispersion of u.

The condition is a retention condition. The corridor is a layer, the expected excess of the signal over the corridor is the premium for the layer the convention declines to write, and the bound is placed where that premium equals the prize it denies to capture, weighted by the probability that a captured revision survives. Two consequences follow for the debate over the width of the symmetric adjustment's corridor. The first is that the question has an empirical answer once the distribution of the unbounded adjustment is known: a proposed width is optimal under the model if and only if the stop-loss premium at that retention matches the revision-survival probability the proposer is prepared to defend, and Section 4.5 computes it from the regime's own series. The second is that the argument from procyclicality for a wider corridor, which is an argument about the truncation term alone, is incomplete without a view on q; a supervisor who thinks captured revisions are rarely reversed should want a narrower corridor than one who thinks the salience of a legislative review protects the bound.

3.4 Auditability

Let x^t denote the published signal history and C the published mapping.

Definition. A convention is *auditable* at t if a third party observing (x^t, C) can reproduce s_t exactly. It is *responsive* if C depends non-trivially on x.

Proposition 3. *A rule is auditable and responsive. Discretion is responsive and not auditable. A frozen parameter is auditable and not responsive. Auditability and responsiveness are jointly*

satisfied only by a rule.

The result is definitional and I claim nothing more for it than that it organizes the cases. Its use is in blocking an inference that is otherwise tempting. Auditability is the property most readily invoked when a professional body asks whether a practice is transparent, and a frozen parameter passes that test perfectly: everyone knows what the number is and everyone can check that it was applied. The proposition makes explicit that passing it is compatible with the standard carrying no information whatever about the state it refers to. Transparency about a number is not the same as transparency about the judgment the number embodies, a distinction the central-banking literature on transparency has drawn in its own terms (Geraats 2002), and only a published mapping delivers the second.

The failure of auditability under discretion deserves a precise statement, because it is not a failure of good faith. From (3), the discretionary standard depends on k, which is observed by the supervisor and, in the relevant detail, not by third parties, and on the weights γ and ϕ, which are not published and in general not articulated even internally. Two identical signal realizations can therefore produce two different standards without any impropriety and without any possibility of an outside party establishing which considerations moved the number.

3.5 Escape clauses and the information a rule cannot carry

A rule written on a signal inherits the signal's omissions. Restore $\tau > 0$: at $t = 2$ the supervisor observes a shift β that no mapping could have anticipated, so that under the rule $s = x$ the standard is wrong by β in every state, at expected cost $\tau^2/2$. Under discretion the supervisor uses β, setting $s_D(x+\beta, k)$ by the analogue of (3), but pays the accommodation cost of Proposition 1. The standard institutional response to this trade-off is an escape clause, and the condition that makes an escape clause compatible with commitment is well understood from the monetary literature (Lohmann 1992; Bernanke and Mishkin 1997): the escape must be costly to invoke, its invocation must be observable, and the occasions for it must be fixed in advance.

Let regime $E(\bar{\beta})$ set $s = x$ unless $|\beta| > \bar{\beta}$, in which case the supervisor sets $s_D(x+\beta, k)$ at invocation cost $\kappa > 0$, with the threshold $\bar{\beta}$ announced at $t = 0$ and the invocation published together with the observed β. The threshold is a pre-specified trigger on a variable exogenous to the firm.

Proposition 4. *(i) Under $E(\bar{\beta})$ the firm's capital $k_E(\bar{\beta})$ tends to k_R as $\bar{\beta} \to \infty$ and to k_D as $\bar{\beta} \to 0$, and it is non-decreasing in $\bar{\beta}$, with $k_D \le k_E \le k_R$, whenever the accommodation discount outweighs the convexity effect of the shift, that is, whenever $\mathbb{E}_\beta[(x+\beta-k)_+] \le (1+\gamma\phi(x))^2\,(x-k)_+$ on the states that matter. (ii) If instead the escape is invoked on a condition that depends on k, and the firm can lower k so as to make invocation more likely, the firm's capital is k_D regardless of the threshold. (iii) Expected supervisory loss under $E(\bar{\beta})$ tends to L_R as $\bar{\beta} \to \infty$ and to $L_D + \kappa$ as $\bar{\beta} \to 0$. If the shift has unbounded support, $L_E(\bar{\beta}) < L_R$ for all sufficiently large finite $\bar{\beta}$: an escape reserved for the largest shifts always improves on the rigid rule. Whether the best threshold also lies below L_D depends on the size of the accommodation loss relative to the information the mapping omits.*

Part (ii) is the operative content and the feature most likely to be lost in implementation. An escape clause conditioned on anything the firm controls – its capital position, its solvency ratio, the state of its balance sheet – reproduces the dependence of (3) and with it the whole of Proposition 1. The firm will again hold less, knowing that holding less makes the escape more likely to be invoked

in its favour. An escape conditioned on a threshold crossed by an exogenous variable cannot be manipulated in this way. The distinction between "we may relax the standard when conditions warrant" and "we may relax the standard when the published indicator crosses a stated threshold, and we will publish that we have done so" is the whole difference between a convention that commits and one that does not, and it is invisible at the level of rhetoric. Part (iii) says when the escape is worth having. The improvement on the rule is a matter of theory; its size is a matter of how much the signal leaves out, and when the signal is nearly sufficient the improvement is negligible. Whether the escape regime beats discretion as well is a quantitative question, answered affirmatively in every calibration of Section 3.7, and the condition is that the accommodation loss of Proposition 1 be larger than what the rule loses by ignoring the shift.

3.6 Endogenous signals and the concentration of the signal population

The signals that proxy for underwriting conditions are loss ratios, combined ratios and reserve development: statistics of the regulated population itself, so that a rule written on them can be influenced by the firms it binds. The symmetric adjustment's equity index is not of this kind; a sector loss-ratio indicator of the sort Section 4.3 requires is. The question is how much that matters, and the answer turns out to be a concentration measure.

Let the signal be a weighted statistic of firm-level observables, $x = \sum_i w_i x_i$ with $\sum_i w_i = 1$, and let firm i be able to lower its own contribution by $m_i \geq 0$ at cost $\frac{a}{2} m_i^2$: through reserving and reporting choices that are themselves actuarial judgments, the reason the reported statistic is manipulable at the margin at all (Petroni 1992; Grace and Leverty 2012). The standard $s = x$ applies to every firm, and firm i's shortfall cost is $\frac{\psi}{2}(s - k_i)_+^2$.

Proposition 5. *In the Nash equilibrium of the manipulation game, firm i lowers its contribution by $m_i^* = \psi\, w_i\, (s - k_i)_+ / a$, and the induced downward displacement of the standard is $(\psi/a) \sum_i w_i^2\, (s - k_i)_+$ to first order. For symmetric firms the displacement is exactly*

$$\bar{x} - s \;=\; \frac{(\psi/a)\, H}{1 + (\psi/a)\, H}\, (\bar{x} - k)_+,$$

where $\bar{x}$ is the unmanipulated statistic and $H = \sum_i w_i^2$ is the Herfindahl index of the signal population. The wedge is increasing in H, vanishes as $H \to 0$, and is largest for a rule written on a single firm's own signal.

The result gives the design criterion its quantitative form. A convention should be written on a population whose concentration is low relative to any one firm's share, or on a market external to the regulated firms altogether. The equity index behind the symmetric adjustment sits at $H \approx 0$ with respect to insurers' influence. A cycle indicator drawn from a fragmented primary motor market sits near it. One drawn from a reinsurance market in which a handful of groups write most of the business does not, and a firm-specific indicator, which is what an internal model's own parameters amount to, sits at $H = 1$. The proposition also says what a supervisor who must use a concentrated population's statistics can do about it: raise a, the cost of moving one's own contribution, which in practice means tying the signal to audited figures with a lag rather than to management estimates.

3.7 Numerical illustration

The propositions are qualitative. A calibrated version of the model, detailed in Appendix B, shows the magnitudes and the shape of the trade-offs. The signal is lognormal with unit mean and a coefficient of variation of 0.30; the shortfall weight is $\gamma = 0.75$ with $\phi(x) = x$; the firm's late-capital cost has $\psi = 3$ against a unit cost of advance capital $c = 0.30$; the non-contractible shift has standard deviation $\tau = 0.15$, that is, 15 per cent of the signal's mean; and the invocation cost is $\kappa = 0.01$, about a sixth of the supervisor's expected loss under the rule. Table 4 collects the parameters and Table 5 the sensitivity analysis.

Figure 1 shows Proposition 1. Under the committed rule the firm holds $k_R = 1.04$, slightly above the mean signal; under discretion it holds $k_D = 0.66$, a reduction of 36 per cent, and the realized standard bends away from the signal above k_D, by 0.14 at a signal of 1.0, by 0.44 at 1.5, and by 0.80 at 2.0 – at twice the mean signal the discretionary standard is 40 per cent below what the signal implies. The firm is under-capitalized relative to the signal in 89 per cent of states under discretion and in 39 per cent under the rule. The supervisor's expected loss is 24 per cent higher under discretion.

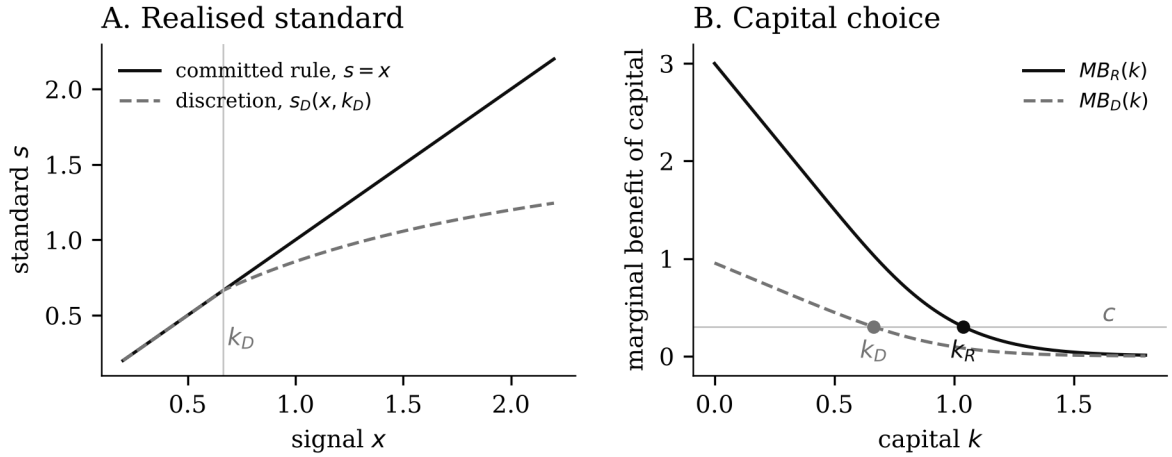


Figure 1: Discretion relaxes the standard where the signal is worst. Panel A: the realized standard as a function of the signal under the committed rule and under discretion, at the discretionary equilibrium capital k_D. Panel B: the firm's marginal benefit of advance capital under each regime; the equilibrium is where it meets the unit cost c. Central calibration of Table 4.

Figure 2 shows Proposition 2 for $b/a = 1$, $q_D = 0.1$ and three values of the reversal probability under a rule revision. The rule yields less distortion than discretion only when the revision interval is shorter than $(1-q_D)/(1-q_R)$ periods: 1.3, 1.8 and 3.0 periods for q_R of 0.3, 0.5 and 0.7. Beyond that the unbounded rule is worse than discretion, and the gap widens linearly with the interval. The bound, here $\bar{\delta} = 2.5$, caps the rule's distortion at $(1 - q_R)\bar{\delta}$ from the third period onward, whatever the interval. Nothing in the figure depends on the calibration of Table 4; the parameters are illustrative, and the qualitative pattern is Proposition 2.

Figure 3 shows Proposition 4. Panel A plots the supervisor's expected loss under the rule with escape as a function of the threshold, in units of τ. At a threshold of zero the escape is invoked in every period and the regime is discretion plus the invocation cost; at a high threshold it is never

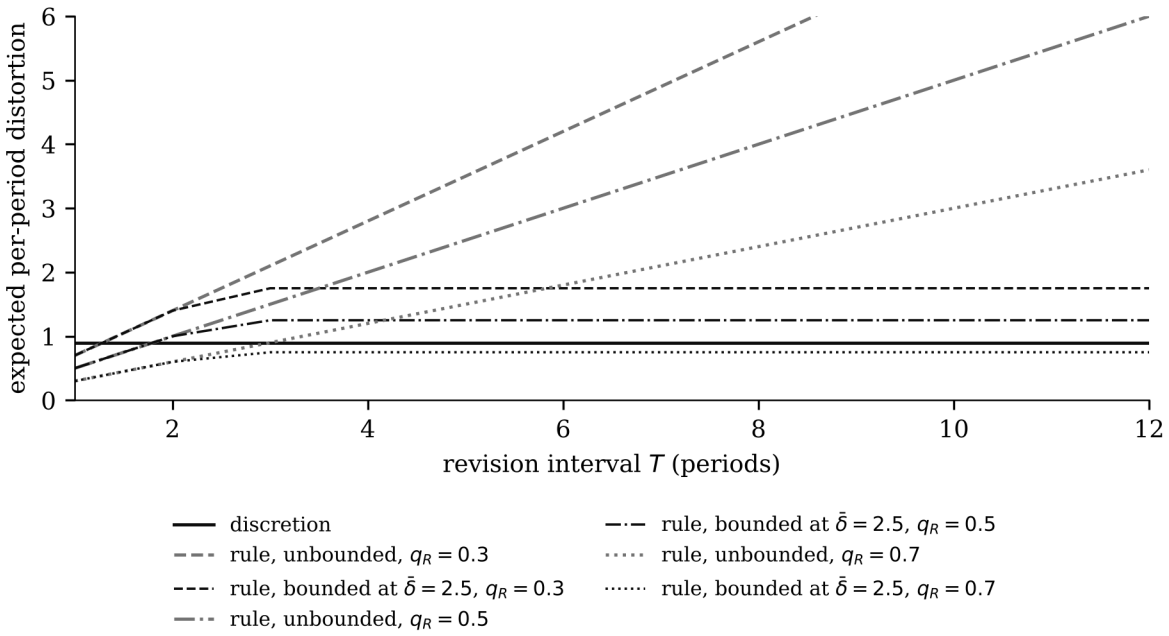


Figure 2: Capture: frequency against persistence. Expected per-period distortion under discretion (horizontal line) and under a rule revised every T periods, unbounded (grey) and with the mapping's output bounded at $\bar{\delta} = 2.5$ (black), for three values of the probability q_R that a captured revision is reversed; $q_D = 0.1$, $b/a = 1$.

invoked and the regime is the rule. In between the loss is U-shaped, with a minimum at a threshold of about 1.35τ, at which the escape is invoked in 18 per cent of periods, the firm holds $k_E = 1.01$, within 3 per cent of k_R, and the loss lies 5 per cent below the rule's and 24 per cent below discretion's. Panel B shows the capital held, rising from k_D to k_R as the threshold rises, which is part (i) of the proposition. The sensitivity rows of Table 5 show part (iii): with $\tau = 0.05$ the escape is never worth invoking and the optimum coincides with the rule, while with $\tau = 0.25$ discretion already beats the rigid rule and the escape beats both by a wider margin.

Figure 4 shows Proposition 5 for $\psi/a = 0.75$ and a shortfall of 0.25: the manipulation wedge rises with the Herfindahl index, nearly linearly over the relevant range, and is negligible for a population of fifty equal firms, a fifth of its maximum for eight, and largest for a rule written on a single firm's own statistic.

In Table 5 the ratio k_D/k_R falls from 0.75 to 0.53 as γ rises from 0.5 to 1.0: the more the supervisor cares about the cost of forced capital-raising, the less capital the firm raises in advance, which is the accommodation mechanism in one line. And the value of the escape clause rises with τ and falls with κ in the way part (iii) describes, with the invocation frequency at the optimum ranging from effectively zero to 37 per cent across the rows; the escape is not a fixed feature of good design but an instrument whose use should scale with how much the signal leaves out.

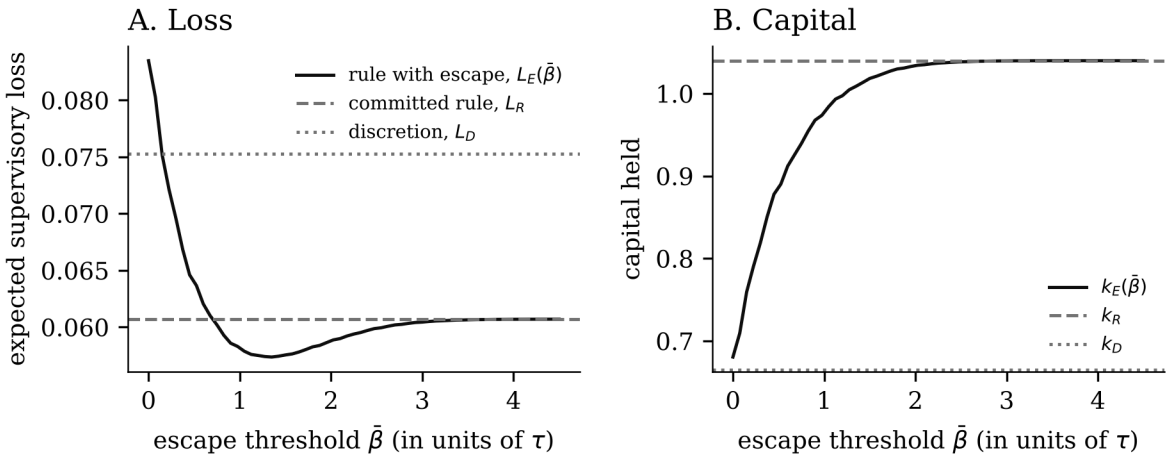


Figure 3: The escape clause. Panel A: expected supervisory loss under a rule with a threshold escape, as a function of the threshold $\bar{\beta}$ in units of τ, against the committed rule and discretion. Panel B: capital held under the same regime. Central calibration of Table 4.

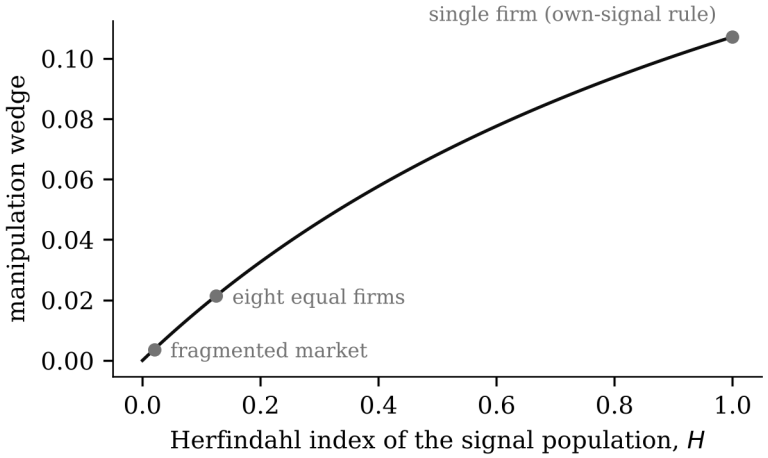


Figure 4: Concentration of the signal population. The downward wedge induced in a rule written on a statistic of the regulated population, as a function of that population's Herfindahl index; $\psi/a = 0.75$, $(s-k)_+ = 0.25$.

4. Three Conventions in the European Solvency Regime

The regime provides three parameters that occupy three cells of the taxonomy of Section 2.3 and illustrate the propositions of Section 3. Table 1 places them; Table 2 summarizes their governance attributes. I take them in order of governance quality, not of importance, then turn in Section 4.4 to what the record of their revision shows, and in Section 4.5 to what the symmetric adjustment's own published series says about the width of its corridor.

Table 1. The three regimes against the three criteria, with the three parameters placed.

	Rule	Discretion	Frozen parameter	Frozen between unannounced revisions
Responsive	yes	yes	no	no
Determinate	yes	no	yes	no at revision
Time-consistent	yes, if the signal is exogenous	no (Proposition 1)	trivially	–
Capture profile	infrequent and persistent; bounded by a cap on output	frequent and transient	at legislative revision only	at revision, with low salience
Auditable	yes	no	yes	partly
Example	symmetric adjustment	a parameter re-set by the supervisor	99.5 per cent standard; cost-of-capital rate	premium risk factors (Annex II)

Table 2. Governance attributes of three Solvency II parameters.

Attribute	Symmetric adjustment	99.5 per cent standard	Premium risk factors
What the number proxies for	the position of equity markets relative to trend	how much safety the regime buys (a preference)	the volatility of segment loss ratios (a state)
Signal	published equity index	none	none
Mapping	published formula, Art. 172 DR	none	none
Bound on output	±10 points; ±13 from 2027	–	–
Revision of the mapping	legislative review, public	ordinary legislative procedure	Commission delegated act after EIOPA advice
Revision occasions	fixed by review clauses	none scheduled	not specified in advance (2015; 2019)
Legal level of the number	Directive Art. 106 (bound); DR Art. 172 (mapping)	Directive Art. 101(3)	DR Annex II

Attribute	Symmetric adjustment	99.5 per cent standard	Premium risk factors
Cell	rule	frozen parameter	frozen between unannounced revisions

4.1 The symmetric adjustment: a rule with a bounded output and a contested bound

The equity capital charge is adjusted by a quantity computed from a published index. The adjustment equals one half of the proportional excess of the current index level over the equal-weighted average of its daily levels across the preceding 36 months, less eight percentage points, and is confined to the interval from -10 to $+10$ percentage points. EIOPA publishes the value monthly. The legal basis is Article 106 of the Directive, which sets the bound, and Article 172 of the Delegated Regulation, which sets the mapping.

Every element of a rule in the sense of Section 2.3 is present. The signal is public and verifiable; the mapping is published in a legal instrument; the output can be computed by a third party as readily as by the supervisor; and the mapping's output is bounded. Against Proposition 3 the convention is auditable and responsive. Against Proposition 1 it is immune to the accommodation channel, because the mapping does not take any undertaking's capital position as an argument: the signal is a market index, exogenous to every individual insurer and, for practical purposes, to the sector. Against Proposition 5 the signal population has a Herfindahl index of effectively zero with respect to the firms the rule binds. The mapping's most severe test to date came in the spring of 2020, when the equity decline took the adjustment to its lower limit and the bound became binding (Aegon Asset Management 2022; EIOPA 2026).

The treatment of the bound is the instructive part, and Section 4.4 sets out the record. Here it is enough to note the architecture. The contest over the parameter takes place at a scheduled, public, documented revision of the mapping, where it attracts argument from identifiable parties with stated reasons; in the intervals between such revisions, the bound caps what any party could gain. The convention concentrates the political economy at the point where it is visible and caps it everywhere else, which is the design Proposition 2 recommends.

The arithmetic of the corridor shows where the bound sits relative to the body of the distribution. Writing d for the proportional deviation of the index from its 36-month average, the adjustment is $\frac{1}{2}(d - 0.08)$, so that the neutral point lies at a deviation of plus 8 per cent and the limits of ± 10 points are reached at deviations of -12 and $+28$ per cent; ± 13 moves those triggers to -18 and $+34$ per cent, and the proposed ± 17 would have moved them to -26 and $+42$ per cent. At the end of March 2020 the deviation stood at -18 per cent, an unbounded adjustment of -13.1 points published as -10. Each widening admits more of the signal's tail and raises the prize available at the next revision by the same amount; Proposition 6 says where the two balance, and Section 4.5 locates the three corridors on EIOPA's own series. The same episode showed that publication frequency is a lever separate from the mapping: EIOPA supplemented the monthly publication with extraordinary updates between March and September 2020 without altering either the formula or the bound.

One qualification. A rule written on a market index inherits that index's properties, including the possibility that the index ceases to represent the exposures it was chosen for. The convention's

responsiveness is only as good as the signal's relevance, and the periodic review of the index composition belongs to the convention and is not a housekeeping matter.

4.2 The 99.5 per cent standard: a number without a convention

The confidence level sits in Article 101(3) of the Directive. There is no signal, no mapping, and no revision trigger; the number changes when the Directive is amended through the ordinary legislative procedure, and the 2025 amending directive, which altered a great deal else, left it alone. In the taxonomy it is a frozen parameter, and by Proposition 3 it is auditable and wholly unresponsive.

The temptation is to read this as the worst case, and I think that reading is wrong, for a reason that yields the paper's main design criterion. A confidence level is not a proxy for a state of the world. It is a statement about how much safety the regime buys, which is to say a preference, and preferences have no business tracking a signal. One would not want the quantity of protection society demands from its insurers to fall automatically when a market index falls; a standard that moved that way would be the procyclicality problem in its purest form. Where the tail beyond the standard lands, whether on policyholders, on guarantee schemes, or on the fiscal capacity of the state, is a political question about who carries catastrophe, not a calibration question (Aslan 2026d). The durability of the 99.5 per cent standard, and its location in a legal instrument that is expensive to amend, are appropriate to what it is.

This gives the criterion that organizes the whole of Section 4:

> Freeze standards that encode preferences; govern by rule standards that proxy for states of the world.

The criterion is not subtle, but it is violated in both directions, and violations in the second direction are the subject of the next subsection. Its corollary is that the first question to ask about any actuarial standard is not how it was calibrated but which of the two kinds of thing it is, a question that frequently goes unasked because the calibration vocabulary treats every parameter as an estimate of something.

A reader will object that the separation is not clean: the 99.5 per cent standard is a preference, but the mapping from that preference to a capital number runs through a loss distribution that is emphatically a state of the world, and the standard formula's modules are where that mapping lives. The objection is correct and it locates the issue. The confidence level is rightly frozen; what should move with the state is the apparatus that converts it into a requirement. The premium risk module is a part of that apparatus that does not.

4.3 The premium risk volatility factors: the least defensible cell

Article 117 of the Delegated Regulation computes the standard deviation for non-life premium risk in each segment as the product of a segment factor set out in Annex II and an adjustment factor for non-proportional reinsurance. The segment factors are fixed numbers: 10 per cent for motor vehicle liability, 8 per cent for fire and other damage to property, 14 per cent for general liability, 17 per cent for each of the three non-proportional reinsurance segments.

These are not preferences. They are estimates of the volatility of a loss ratio, which is to say proxies for a state of the world, and the state they proxy for is not constant. The underwriting cycle is among

the most durable regularities in non-life insurance, documented internationally since Cummins and Outreville (1987) and surviving every subsequent re-examination; combined-ratio series exhibit regime structure rather than stationary variation around a constant mean, with volatility differing between hard and soft states. In work developed separately I estimate a two-state Markov-switching model on United States property-casualty combined ratios and on the Lloyd's market and find a regime spread that the standard formula's time-invariant factors cannot accommodate, with the consequence that required capital moves against conditional underwriting risk rather than with it (Aslan 2026a). The point relevant here is not the calibration but its institutional form: a quantity that varies with an observable state is represented by a parameter that does not vary at all.

By the criterion of Section 4.2 this parameter belongs in the rule cell. The remedy has the shape of the symmetric adjustment, a published mapping from a published cycle indicator to an adjustment of the volatility factor, bounded by a corridor, and that is the proposal developed in calibration terms in Aslan (2026a). What the present framework adds is the reason the proposal should take that form rather than another, and the reason is Proposition 1. The alternative to a rule here is not the status quo; it is supervisory discretion over the factor, which has been advocated under the heading of dynamic or judgment-based calibration. Proposition 1 says what that would produce. A factor set at the supervisor's discretion, after undertakings have set their capital, under a loss function that correctly includes the cost of forcing capital-raising into a hard market, is a factor that will be set lowest when the cycle indicator is highest. The procyclicality that a dynamic calibration is introduced to cure is the procyclicality that discretionary dynamic calibration reproduces. Proposition 5 adds the constraint on the signal: a cycle indicator for a fragmented primary market can be drawn from that market's own statistics with little manipulation risk, whereas one for the reinsurance segments, where a few groups write most of the business, should be drawn from a population wider than the firms it binds, or from audited figures with a lag.

The current arrangement is also not what it appears. The Annex II factors are not frozen in the sense of Section 4.2; they are revised, at intervals, through EIOPA recalibration exercises whose methodology is published but whose occasions are not specified in advance. In the taxonomy this is neither a rule nor a frozen parameter proper. Between revisions the standard is unresponsive; at revision it is set by a process whose timing was not announced, which is the definition of indeterminacy. The parameter thus inherits the informational inertness of the frozen cell and the unauditability of the discretionary one.

4.4 What the record shows, 2019–2027

The model makes three predictions about how parameters in the three cells should behave over time, and the regime's recent history allows each to be checked against the record.

The first concerns the frozen-between-revisions cell. If the Annex II factors are revised at unannounced occasions by a technical process, revisions should be rare, partial, and clustered at whatever review the Commission happens to schedule. That is what happened. The factors adopted in 2015 were revised once, by Delegated Regulation (EU) 2019/981 following EIOPA's advice on the review of the standard formula, and the revision changed three of the twelve segments: credit and suretyship moved from 12 to 19 per cent for premium risk and from 19 to 17.2 per cent for reserve risk; legal expenses from 7 to 8.3 and from 12 to 5.5; assistance from 9 to 6.4 and from 20 to 22. The motor liability, property, general liability and non-proportional reinsurance

factors, which carry most of the capital, were not touched, and the 2026 amending regulation left the segment factors unchanged, confining its changes to the module to the treatment of non-proportional and adverse development reinsurance. A decade of underwriting cycle, including the hardest casualty market in two decades, left the liability factors where they were. That is the informational inertness the taxonomy predicts, and it is not a criticism of the recalibration work, which did what a recalibration can do; it is a criticism of the cell.

The second concerns the rule cell and specifically Proposition 2. If the bound on a rule's output is where the political economy concentrates, then a scheduled revision of the mapping should produce a visible contest over the bound, among identifiable parties with stated reasons, and an outcome that reflects it. The 2020 review produced exactly that. EIOPA's opinion of December 2020 proposed widening the corridor from ± 10 to ± 17 percentage points (EIOPA 2020). The Commission adopted the proposal in its legislative text of September 2021 (European Commission 2021). National industry bodies opposed it on the record: Insurance Sweden, whose members hold the largest equity allocations in Europe, argued in March 2021 that a wider corridor would raise the volatility of capital requirements and could move solvency ratios perversely in large equity declines (Insurance Sweden 2021). The co-legislators settled on ± 13 in Directive (EU) 2025/2, and Delegated Regulation (EU) 2026/269 carried the figure into Article 172 with effect from 30 January 2027. Whether thirteen is the right number is not the point; the point is that the contest was dated, documented, attributable, and bounded. Since the regime took effect in 2016 nothing anyone did could move the symmetric adjustment outside ten points, and the one occasion on which the bound moved was the most visible decision point the regime possesses. That is the salience channel of Proposition 2 doing what it is supposed to do. The same review also shows the recursion of Section 3.3: the bound sits in the Directive and the mapping in the Delegated Regulation, so moving the bound required the ordinary legislative procedure while the mapping followed by delegated act, and that is why the contest took place where it did.

The third concerns the frozen cell. A frozen parameter is revised only at legislative review, and when it is, the revision should be a one-shot movement with the whole prize at stake. The cost-of-capital rate behind the risk margin was 6 per cent from the regime's inception. The 2020 review reduced it to 4.75 per cent and introduced a time-dependent factor that lowers the risk margin further for long-dated business, a reduction in the risk margin of roughly a fifth across the sector (Actuarial Association of Europe 2025). I draw no conclusion about whether 4.75 is better than 6; the model does not say that frozen parameters are revised in the wrong direction, only that they are revised rarely and wholesale, and that the salience of a legislative review is what the regulated sector's effort is concentrated on. The pattern of the 2020 review, a decade of stability and then simultaneous movement of several frozen parameters at one visible occasion, is the pattern Proposition 2 describes.

None of this is a test in the econometric sense; three parameters and one review cycle cannot be. What the record does is show that the objects the model describes exist, behave as the taxonomy says they should, and can be distinguished from one another by the attributes in Table 2. Section 6.3 says what a test would require.

4.5 The corridor in the data, 1991–2026

The documentary record of Section 4.4 shows the objects of the model existing and behaving as the taxonomy says they should. The symmetric adjustment allows one step more, because its owner publishes not only the standard but the whole convention: the index, the 36-month average, the deviation, the unbounded output of the mapping and its truncation, for every calendar day. EIOPA's workbook for end-September 2026 carries the series from 30 June 1991, extending its index back a quarter of a century before the regime took effect (EIOPA 2026). The published monthly value is the month-end entry, and the bounded series equals the truncation of the unbounded one at ±10 points exactly, which confirms that the mapping of Article 172 is the mapping in use. Two samples are used below: the 129 month-ends of the regime period, January 2016 to September 2026, and the 424 month-ends of the full series. The second includes the two equity collapses that the corridor was designed with in view; the first is the record the corridor has actually governed. Table 6 summarizes both, and Figure 5 shows the unbounded adjustment against the three corridors that have been in force, legislated, or proposed.

Table 6. The unbounded symmetric adjustment, EIOPA series, percentage points.

	Regime period, month-ends	Regime period, days	Full series, month-ends	Full series, days
Observations	129	3,926	424	12,877
Mean	0.67	0.69	0.62	0.64
Standard deviation	4.84	4.87	9.05	9.14
Minimum	−13.07 (31 Mar 2020)	−18.01 (18 Mar 2020)	−26.28	−28.04 (9 Mar 2009)
Maximum	10.35 (28 Feb 2026)	10.69 (7 Aug 2026)	26.87	29.92 (6 Apr 1998)
Share beyond ±10	0.023	0.024	0.224	0.234
Share beyond ±13	0.008	0.006	0.134	0.136
Share beyond ±17	0.000	0.001	0.080	0.085
Stop-loss premium at 10	0.029	0.050	1.324	1.376
Stop-loss premium at 13	0.001	0.015	0.790	0.830
Stop-loss premium at 17	0.000	0.001	0.385	0.409

Source: EIOPA (2026), sheet "Calculations"; author's calculations. The stop-loss premium at $\bar{\delta}$ is $\mathbb{E}[(|u| - \bar{\delta})_+]$ in percentage points.

The record has three features that bear on the argument. The first is how rarely the bound has bound. In the regime period the published value lay outside ±10 at three of 129 month-ends: March 2020 at −13.07, April 2020 at −10.26, and February 2026 at +10.35, the first and so far only occasion on which the cap rather than the floor has engaged. On a daily basis the unbounded value lay outside the corridor on 95 of 3,926 days, 75 of them in 2020, when it reached −18.01 on 18 March. A ±13 corridor would have bound at one month-end, March 2020, by seven hundredths of a point; a ±17 corridor at none. The full series tells a different story: outside ±10 on 22 per cent of month-ends, below −10 for 777 consecutive days from August 2001 to October 2003 and for 410 days from September 2008 to October 2009, above +10 for 600 consecutive days from January

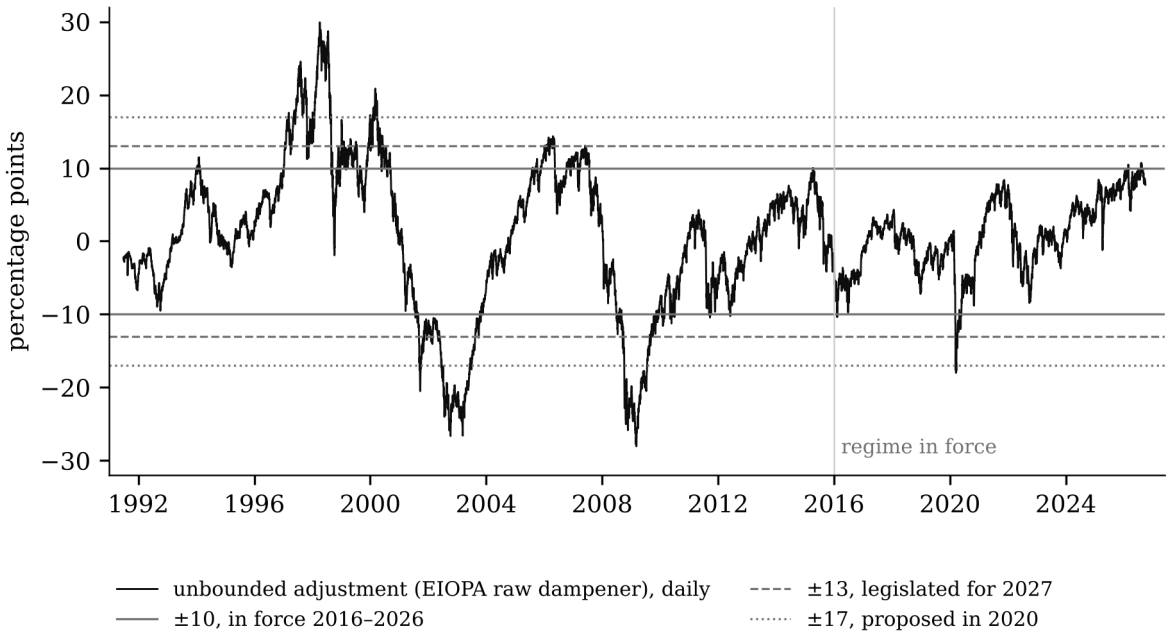


Figure 5: The unbounded symmetric adjustment, 30 June 1991 to 30 September 2026, daily, against the corridor in force since 2016 (±10), the corridor legislated for 2027 (±13) and the corridor proposed in 2020 (±17). Source: EIOPA (2026); the series before 2016 is EIOPA's extension of its index.

1997 to August 1998. The bound was set in the Directive of 2009 and the mapping in the delegated regulation of 2014, on a history in which the corridor would have bound nearly a quarter of the time. It has since governed a decade in which it bound one month in forty.

The second concerns the two arguments made in the 2020 review, and Table 7 puts numbers to both. Insurance Sweden's objection was volatility: the standard deviation of the monthly change in the equity capital charge, 39 points plus the adjustment, would have been 2.01 points under ±13 against 1.91 under ±10 over the regime period, five per cent higher, and 1.94 against 1.80 over the full series, eight per cent higher; under ±17 the full-series figure is 2.08, sixteen per cent higher, and without any bound 2.35. The objection is correct in sign and modest in size. EIOPA's argument was relief: at end-March 2020 the charge was published at 29.0 and would have been 26.0 under ±13 and 25.9 under ±17, three points of additional relief at the trough and almost nothing beyond thirteen; at end-December 2008 the three corridors give 29.0, 26.0 and 22.0. The relief a wider corridor buys is concentrated in collapses of the 2008 kind, and the 2020 episode, severe as it was on a daily basis, did not reach the region in which the third corridor adds to the second. The two arguments are both right about their own term of (6) and silent about the other, which is the position Proposition 6 was written to resolve.

Table 7. Counterfactual corridors: the equity capital charge under alternative bounds.

	±10	±13	±17	no bound
Regime period, 129 month-ends				
Months at the bound	3	1	0	–

	±10	±13	±17	no bound
S.d. of monthly change in the charge	1.91	2.01	2.01	2.01
Largest monthly fall	−5.02	−7.06	−7.13	−7.13
Lowest charge (end-March 2020)	29.00	26.00	25.93	25.93
Highest charge	49.00	49.35	49.35	49.35
Stop-loss premium at the bound	0.029	0.001	0.000	–
Revision-survival probability that rationalizes the bound, q	0.997	1.000	1.000	–
90 per cent block-bootstrap interval for q	0.992–1.000	1.000–1.000	1.000–1.000	–
Full series, 424 month-ends				
Months at the bound	95	57	34	–
S.d. of monthly change in the charge	1.80	1.94	2.08	2.35
Largest monthly fall	−7.51	−7.51	−7.51	−13.64
Lowest charge	29.00	26.00	22.00	12.72
Highest charge	49.00	52.00	56.00	65.87
Stop-loss premium at the bound	1.324	0.790	0.385	–
Revision-survival probability that rationalizes the bound, q	0.868	0.939	0.977	–
90 per cent block-bootstrap interval for q	0.776–0.952	0.893–0.981	0.959–0.993	–

Source: EIOPA (2026); author's calculations. The charge is 39 points plus the bounded adjustment. The rationalizing q is $1 - \mathbb{E}[(|u| - \bar{\delta})_+]/\bar{\delta}$, the value at which the bound satisfies the condition of Proposition 6; the intervals are from a circular block bootstrap with 36-month blocks and 5,000 replications (Appendix B).

The third is what Proposition 6 says about the widths themselves, and Figure 6 shows it. The condition of the proposition is that the stop-loss premium of the unbounded adjustment at the bound equal $(1 - q)$ times the bound, so each candidate width can be read as the revision-survival probability that would make it optimal. On the full series the ±10 corridor is optimal if $q = 0.87$, that is, if seven captured revisions of the bound in eight are reversed; ±13 requires $q = 0.94$; ±17 requires $q = 0.98$. Each widening is, under the model, a claim that the legislative review protects the bound more reliably than the previous width assumed. The estimates carry the uncertainty one expects of a series built on a 36-month window, in which 424 month-ends amount to a handful of independent episodes: a block bootstrap with three-year blocks puts the q that rationalizes ±10 between 0.78 and 0.95, ±13 between 0.89 and 0.98, and ±17 between 0.96 and 0.99, and the optimal width at $q = 0.95$ between 10 and 16 points. The data therefore do not separate the corridor in force from the corridor legislated for 2027 at any q above 0.9; what they do establish is the ordering, and that ±17 is rationalized only by a revision-survival probability above 0.96.[5] On the regime

[5]The corridor is symmetric by construction but the distribution is not: the neutral point of the mapping sits at a deviation of plus 8 per cent, and on the full series the lower tail accounts for 0.74 of the 1.32 points of stop-loss premium at 10 and the upper tail for 0.58. Proposition 6 applies side by side to separately set bounds, with the lower bound governed by $\mathbb{E}[(-u - \bar{\delta}^-)_+] = (1 - q^-)\bar{\delta}^-$ and the upper by its mirror image; at 10 points the two sides are rationalized by $q^- = 0.93$ and $q^+ = 0.94$, so the symmetric design costs little, but the two sides answer to different

period alone the stop-loss premium at 10 points is three hundredths of a point and the corridor is rationalized only by $q = 0.997$: for values of q between 0.5 and 0.95 the optimal width lies between 3 and 7 points, and every corridor that has been in force, legislated or proposed is wider than the optimum, because the tail the corridor exists to truncate has scarcely appeared since 2016. The dispute of 2020 to 2024 was conducted over a region of the distribution that the regime's own decade had visited on three month-ends.

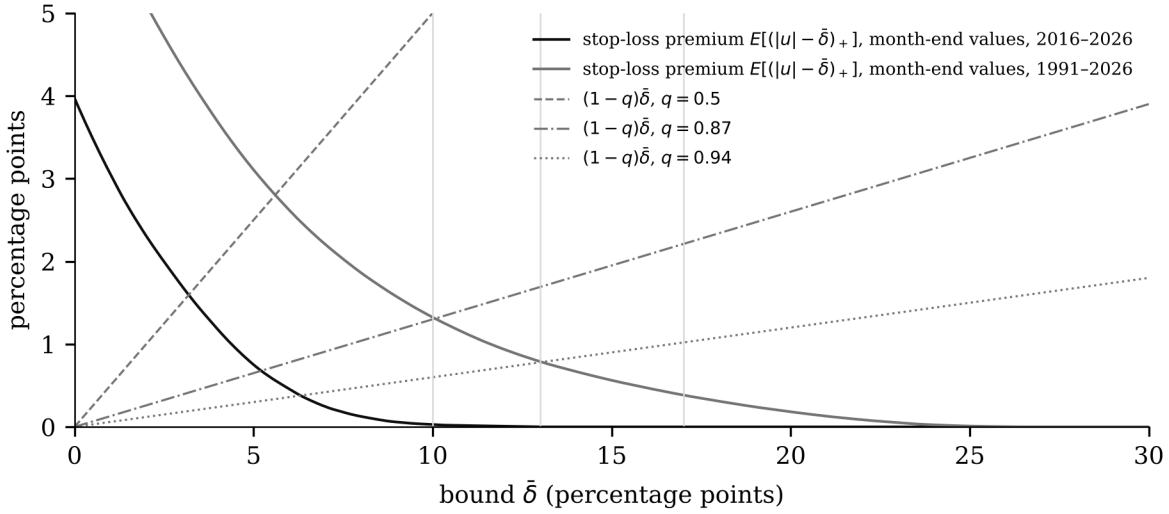


Figure 6: The optimal bound (Proposition 6). The stop-loss premium of the unbounded adjustment, $\mathbb{E}[(|u| - \bar{\delta})_+]$, as a function of the bound, for the regime period (black) and the full series (grey), against $(1-q)\bar{\delta}$ for three values of the revision-survival probability. The optimal bound for a given q is where the curve for the relevant sample meets the line; the vertical lines mark 10, 13 and 17 points. Source: EIOPA (2026); author's calculations.

The point of the exercise is not that the corridor should be narrowed to five points. It is that the choice between the two samples is itself a convention about the tail, of the kind Section 4.2 discussed under the 99.5 per cent standard: whether 1997 to 2003 and 2008 to 2009 belong to the distribution the corridor is written on is a judgment no measurement settles, and the width that follows moves from five points to fourteen with the answer. That is the sense in which the corridor cannot be calibrated and must be governed. A calibration returns a number conditional on a sample; the convention is the choice of sample, of the loss function that weighs truncation against capture, and of the probability that the bound survives the next review, and those three choices, made explicitly and published, are what Proposition 6 asks for. Two qualifications attach. The series before 2016 is EIOPA's reconstruction with its present index, not a record of published standards; and the exercise identifies the q that would rationalize each width under the loss function (6), not the q anyone holds. What the data settle is narrower and still useful: on the regime's own decade the bound has cost almost nothing in transmitted information, on the long history it would have cost a great deal, and the two widths now on the statute book sit between the optimum for the one sample and the optimum for the other.

interests, since the industry's pressure runs on the floor and the supervisor's on the cap, and nothing in the model requires the two q's to coincide.

5. Unvalidatable Inputs and the Certification Problem

5.1 The structure of the problem

The preceding sections concern parameters set by supervisors. The same structure now arises at the level of the individual actuary, and the mechanism that brings it about is the entry of statistical learning methods into pricing, reserving, and underwriting (Richman 2021). Consider an actuary certifying a rate or a technical provision one of whose inputs is produced by a model the actuary did not build, cannot inspect at the level of its parameters, and in many cases has licensed from a vendor under terms that preclude inspection. The number is produced; the actuary signs. The question of what is being certified has an uncomfortable answer. It is not that the number is correct, because no evidence available to the signatory establishes that. It is not that the method is standard, because the method is novel by construction. In the framework of this paper the position is exact: the signal is not public, the mapping is not published, and the convention is accordingly neither auditable nor, from outside, responsive in any checkable way.

By Proposition 3 this is the discretionary cell. In one respect it is worse than supervisory discretion, which at least has an identifiable decision-maker who can be asked for reasons and held to them. A model output has no owner in that sense: the vendor disclaims, the actuary cannot reconstruct, and the question "why this number rather than another" has no addressee. In another respect it is better, in that the accommodation mechanism of Proposition 1 requires an interested party with a loss function, and a model has none. That asymmetry matters for the appropriate response. The problem here is not time inconsistency; it is auditability, and the remedies differ accordingly. The model-risk literature has reached a related conclusion from the other side, finding that the dispersion across risk models of the numbers they produce for the same portfolio is itself large and unstable (Danielsson et al. 2016); what that literature treats as a property of models, this paper treats as a property of the conventions that govern them.

5.2 The three criteria applied

Time consistency is engaged where the model is retrained or re-selected after the business has been written. A model whose specification is revisited whenever its output becomes inconvenient is a discretionary standard in the sense of Section 3.2, with the firm's own position doing the work that k does in (3). The governance response is the one Proposition 4 suggests: fix the retraining schedule in advance, specify the triggers for off-schedule retraining on performance statistics the business unit does not control, and record invocations.

Capture is engaged wherever the model's output bears on a commercial interest within the firm, which in pricing is always. The relevant distortion is not corruption but the gradual accumulation of specification choices that lean one way, each defensible alone. Proposition 2 indicates the remedy that does not depend on detecting any individual choice: bound the output. A corridor on the permitted departure of a model-driven rate or reserve from a documented benchmark caps the cumulative distortion regardless of how the model arrived at its number and regardless of who influenced it, and does so without requiring anyone to adjudicate the specification choices one at a time.

Auditability is the binding constraint. What can be published is not the parameters, which would inform no one, but the mapping at the level at which the convention operates: which inputs the

model consumes, what the published benchmark is, what the corridor is, what the monitoring statistics are, and on what triggers the model is revised. That is a convention in the sense of Section 2.1, and it is auditable in the sense of Section 3.4, notwithstanding that the function inside remains opaque. The distinction between auditing a model and auditing the convention that governs it is the practically useful one, because the first is frequently impossible and the second is not. EIOPA's governance principles for artificial intelligence in insurance (EIOPA 2021) ask for explainability proportionate to the use; the present framework says what proportionate explainability consists of when the model itself cannot be explained.

5.3 A minimum governance disclosure

Table 3 states the elements a disclosure would have to contain for the convention governing an unvalidatable input to be auditable in the sense of Section 3.4. Each element corresponds to a proposition.

Table 3. Minimum disclosure for the convention governing an unvalidatable input.

Element	What must be published	Proposition
Inputs	the list of variables the model consumes, and their sources	3
Benchmark	a documented reference rate or reserve, computable by a third party	3
Corridor	the maximum permitted departure of the model output from the benchmark	2
Monitoring statistics	the performance measures tracked, defined so that the business unit cannot move them	1, 5
Revision schedule	the dates at which the model will be re-estimated or re-selected	1, 4
Escape triggers	the thresholds on the monitoring statistics at which off-schedule revision is permitted, and the obligation to record each invocation	4
Owner	the function that owns the convention and signs the record	3

None of the seven elements requires access to the model's internals, which is the point. A reviewer holding the disclosure can verify that the output lay within the corridor, that revisions occurred on schedule or on a recorded trigger, and that the monitoring statistics were those announced. What

the reviewer cannot verify is whether the number is right, and the disclosure does not pretend otherwise.

An illustration gives the scale. A motor insurer prices with a gradient-boosted model whose relativities no actuary can reproduce from the fitted object. The convention governing it would publish the input list; the benchmark, a generalized linear model on the same data whose rate any reviewer can recompute; a corridor, say 15 per cent either side of the benchmark rate at the level of the rating cell, beyond which the boosted output is not applied; the monitoring statistics, out-of-sample loss ratio and calibration by cell, computed by a function that does not report to the pricing unit; an annual re-estimation date; an escape trigger, two consecutive quarters with the monitored loss ratio outside a stated band, with every off-schedule revision logged; and the signature of the actuary who owns the record. The boosted model is as opaque after this as before. The number it produces is now the output of a convention that can be audited in the sense of Section 3.4, and the illustrative figures are the kind of parameters an owner would have to defend.

The disclosure also answers a demand the law now makes. Regulation (EU) 2024/1689 lists systems used for risk assessment and pricing of natural persons in life and health insurance among its high-risk uses, and attaches to them a risk management system, record-keeping, transparency towards those who deploy the system, and human oversight.[6] Those obligations are written in the vocabulary of systems; what the present framework adds is what they amount to when the system cannot be explained. Explainability of a function that resists explanation is not deliverable; publication of the convention that governs the function is, and the seven elements are a specification of it. The same reading applies to the professional standards that already require actuaries to document work so that another actuary can assess it: the seven elements say what that documentation must contain when the model inside it cannot be reconstructed.

5.4 What the actuary certifies

The professional question this analysis answers is narrow but consequential. When an actuary certifies a number whose basis cannot be reconstructed, the certification cannot be a claim about the number's accuracy, and a profession that allows it to be read as one is exposed to a demand it cannot meet. The claim the evidence does support is a claim about governance: that the number was produced by a stated mapping from stated inputs, that the mapping was fixed in advance of the period it governs, that its output is bounded, that the occasions for its revision were specified before they arose, and that the record of all of this is available for inspection.

That is a weaker claim than accuracy, and it is the stronger professional position, because it concerns something the actuary controls and can demonstrate. It also puts the ethical content of the work where the technical skill already is: in constructing and governing the convention, which means choosing the signal, writing the mapping, setting the bound and specifying the triggers, and then in asserting no more for the number than its provenance will bear. Where skill and restraint coincide in this way, which is where ethics meets expertise in the colloquium's phrase and the condition this paper has tried to characterize, the profession is doing something it can defend; where they come apart, which is what happens when an unvalidatable output is certified as though it were an estimate, the defence is unavailable and the exposure is real.

The same structure, at the opposite end of the severity scale, governs the probabilities now

[6] Regulation (EU) 2024/1689 (Artificial Intelligence Act), Annex III, point 5(c); Articles 9, 12, 13 and 14.

demanded of the profession in connection with catastrophic and existential technological risk, where no measurement technology delivers the quantity at all and the only defensible institutional treatment is to govern the number rather than to estimate it (Aslan 2026b, sec. 4.5). The two cases look very different and are the same case. In both, the demand is for a number the evidence cannot supply; in both, the appropriate response is not silence but a published, owned, bounded, trigger-revised convention; and in both, the professional claim attaches to the convention and not to the number.

6. Discussion

6.1 Relation to the literatures drawn on

The structure of Section 3.2 is that of Kydland and Prescott (1977) and Barro and Gordon (1983), and I claim no novelty for it. What differs is the object. In the monetary setting the inconsistency lies between an announced and a realized action; here the authority chooses a published parameter that defines what compliance is, so the inconsistency is invisible in the compliance record. Mechanisms that work by exposing a gap between what was promised and what was done, which is most accountability mechanisms, have nothing to expose, because after a re-parameterization the standard and the conduct agree by construction. What remains is publication of the mapping in advance, and that is why auditability carries more weight here than in the monetary setting.

Proposition 2 concerns the timing of influence, not its magnitude. The specific result that durability cuts against rules under capture while favouring them against time inconsistency is, so far as I can establish, not in the regulatory-design literature, and it generates the one unambiguous recommendation the paper makes: bound the mapping's output. A bound is the only instrument considered here whose performance does not depend on the political economy of the revision process.

That solvency regimes can amplify cycles is well established (Danielsson et al. 2001; Plantin, Sapra and Shin 2008; Adrian and Shin 2010; European Systemic Risk Board 2015). The contribution here is to identify a procyclicality channel that runs through the governance of the parameter rather than through its level: even a correctly calibrated parameter, if revised at discretion, will be revised in the amplifying direction, for reasons that have nothing to do with the calibration and cannot be fixed by improving it.

6.2 Two objections

A reader may grant everything above and observe that the recommendation reduces to indexing parameters to observable indicators, which supervisors do already and which requires no apparatus. The observation is fair as to the responsive element and misses the other two. Indexation alone reproduces the capture problem at the point where the index and the mapping are chosen, and it provides no answer to the question of what happens when the index fails. The design this paper arrives at is a triple – responsive mapping, bounded output, pre-specified escape on an exogenous trigger – and the propositions say why each element matters: the mapping against Proposition 1, the bound against Proposition 2, the escape against the brittleness Proposition 4 addresses, with Proposition 5 constraining the population the mapping may be written on. A regime with any two of the three elements is defective in a way the third repairs.

A second reader may hear the conventions reading as a licence to stop estimating. It is not, and

the scope condition of Section 2.1 is where the line falls. Where a quantity is estimable, estimate it; the argument applies only to quantities that are not, and the practical difficulty is that the regime does not distinguish the two cases in its own vocabulary. A standard formula parameter and an estimated claims frequency are both called calibrations and both carry decimal places. The first contribution of the conventions reading is to force the distinction into the open, and the second is to supply a governance apparatus for the side of it that has none.

6.3 Scope, limitations, and what would revise the conclusions

The model is a single-signal, quadratic-loss environment with one representative firm, and three of its features are restrictive in ways worth naming. It takes the supervisor's loss function as given and benevolent, so the capture analysis of Section 3.3 is appended to the model and not derived within it. It treats the influence technology of Section 3.3 and the manipulation technology of Section 3.6 as separate, when in practice the sector that lobbies over a mapping is the sector whose statistics feed it. And it says nothing about the political economy of convention-setting beyond treating salience as a parameter, when the question of which forums make a revision salient is what a design recommendation needs answered. Section 5 applies the framework; its claims about certification follow from the taxonomy, not from the propositions. The numerical results of Section 3.7 are an illustration of magnitudes under one calibration, not an estimate of anything.

Several developments would bear directly on the conclusions. Evidence that supervisory parameter revisions are uncorrelated with the state of the regulated sector's balance sheets would tell against the empirical relevance of Proposition 1; the record of recalibration exercises across jurisdictions, dated against sector solvency ratios, is where that could be tested, and Section 4.4 is the documentary precursor of such a test rather than a substitute for it. Evidence that bounded mappings have been widened in episodes of sector stress, rather than at scheduled reviews, would tell against the bound's efficacy and in favour of a more pessimistic reading in which every element of the convention is captured together; the 2020 review is one observation against that reading, not a refutation of it. The calculations of Section 4.5 depend on the loss function (6), which weighs truncation and capture in the same prudential units; a loss function that weighed the relief delivered in a collapse more heavily than the symmetric error would move the optimal widths outward, and the exercise should be read as locating the corridors on the trade-off, not as a recommendation of a width. And a demonstration that model-governance disclosure of the kind Table 3 proposes is reconstructible by third parties in practice, and not only in principle, is a precondition for the paper's recommendation to be operational as well as coherent.

7. Conclusion

The profession certifies numbers that cannot be estimated, and has done so for as long as there has been solvency regulation. Better data will not correct the practice, because it is the normal condition of standard-setting at the tail, and it is becoming the normal condition of ordinary actuarial work as methods whose internals cannot be inspected move into pricing and reserving.

Treating such a number as a convention – published, owned, revised by stated rule – moves the question from one that has no answer to one that institutions have answered before. Discretionary revision relaxes the standard where the signal is worst and leaves no record of having done so, which is why time consistency is a governance property and not a calibration property. Rules trade the

frequency of capture against its persistence, so that the bound on a mapping's output, and not the mapping, limits distortion, and the bound has an optimal width where the expected excess of the signal over the corridor equals the prize the corridor denies to capture. An escape clause earns its place only when the signal leaves much out and its trigger is one the regulated firm cannot move, and a rule written on the regulated population's own statistics is manipulable in the measure of that population's concentration. The architecture that follows – published mapping, bounded output, escape on a pre-announced exogenous trigger, signal drawn from a dispersed or external population – already exists in the European regime, in the symmetric adjustment, alongside instruments built otherwise. The 2020 review, in which the one bound in the regime was contested in public and moved by a legislative act while the liability volatility factors passed through a full underwriting cycle untouched, is the taxonomy's prediction in the regime's own documents; and the regime's own series shows the bound engaging at three month-ends in a decade, costing almost nothing in transmitted information over that decade and a great deal over the two collapses before it, with the widths of 10, 13 and 17 points rationalized on the long history by revision-survival probabilities of 0.87, 0.94 and 0.98. Which history the corridor is written on is a judgment about the tail, not a measurement, and that is the paper's thesis in a single parameter.

The criterion for sorting standards is to freeze those that encode preferences and to govern by rule those that proxy for states of the world. The confidence level belongs to the first class and is correctly frozen; the premium risk volatility factors belong to the second and are not governed at all, being neither responsive between revisions nor determinate at revision.

For the individual actuary the consequence is a reframing of what a signature asserts. A certification of an unvalidatable number cannot be a claim about accuracy and should not be read as one. It can be a claim that the number was produced by a stated mapping from stated inputs, fixed in advance, bounded in output, revised on announced triggers, and recorded throughout: a claim about the governance of a convention, weaker than accuracy and the strongest thing the evidence supports. The expertise is in building the convention well. The ethics is in not claiming for the number more than its provenance will carry.

Three extensions follow directly. A model in which the sector that lobbies over the mapping is also the sector whose statistics feed it would join Sections 3.3 and 3.6. An empirical study of parameter revisions across supervisory jurisdictions, testing whether revisions correlate with the state of the sector's balance sheet, would put Proposition 1 at risk in the way a proposition should be. And a field test of the disclosure of Section 5.3, establishing whether third parties can in practice reconstruct a model-driven number from it, would turn a recommendation into a standard, which is, after all, what the argument says such things should become.

Appendix A: Proofs

Proof of Proposition 1. Under discretion the supervisor at $t = 2$ minimizes (2) in s with k fixed and $\mathbb{E}[\theta \mid x] = x$. The objective is strictly convex in s, and for an interior solution with $s > k$ the first-order condition $(s - x) + \gamma\phi(x)(s - k) = 0$ yields (3); where $x \leq k$ the shortfall term is inactive and $s = x$. Expression (4) follows by substitution, and $\partial(x - s_D)/\partial(\gamma\phi) = (x - k)/(1 + \gamma\phi)^2 > 0$ for $k < x$.

For the capital comparison, write the firm's marginal benefit of capital under discretion

as $MB_D(k) = \psi\,\mathbb{E}[(x-k)_+/(1+\gamma\phi(x))^2]$ and under the committed rule $r(x) = x$ as $MB_R(k) = \psi\,\mathbb{E}[(x-k)_+]$. Since $(1+\gamma\phi(x))^2 > 1$ wherever $\gamma\phi(x) > 0$, and the two integrands agree elsewhere, $MB_D(k) < MB_R(k)$ for every k with $\Pr(x > k) > 0$. Both functions are continuous and strictly decreasing in k on the relevant range, and each optimum solves $MB_i(k) = c$ (or is the corner $k = 0$ if $MB_i(0) \le c$, in which case the ordering is weak). Hence $k_D < k_R$ whenever the rule's optimum is interior, and $s_D(x, k_D) < x$ whenever $k_D < x$ by (4). ■

Proof of Proposition 2. Under discretion the interested party solves $\max_{\delta\in[0,\bar\delta]} b\delta - \frac{a}{2}\delta^2$ each period; the unconstrained optimum is b/a and the constrained optimum $\min\{b/a, \bar\delta\}$. Under a rule revised every T periods the same party solves $\max_{\delta\in[0,\bar\delta]} Tb\delta - \frac{a}{2}\delta^2$ once per T periods, with optimum $\min\{Tb/a, \bar\delta\}$, and the resulting distortion applies in each of the T periods. Multiplying by the probability that the distortion survives contest gives the stated per-period expectations. In the unbounded case the rule yields the lower distortion if and only if $(1-q_R)Tb/a < (1-q_D)b/a$, that is, if and only if $1 - q_R < (1-q_D)/T$, which rearranges to the condition in the statement. Where $\bar\delta \le b/a$ the bound binds under both regimes and the distortions are $(1-q_D)\bar\delta$ and $(1-q_R)\bar\delta$. ■

Proof of Proposition 3. Under a rule, $s_t = C(x^t)$ with C and x^t published, so s_t is reproducible, and C depends non-trivially on x by the definition of the regime. Under discretion, s_t is given by (3) and depends on k and on (γ, ϕ), none of which is in (x^t, C); two states with identical x^t and distinct k generate distinct s_t, so no function of (x^t, C) reproduces s_t. Under a frozen parameter $s_t = \bar s$ is reproducible and independent of x^t, so responsiveness fails. The three cases exhaust the occupied cells of the classification in Section 2.3. ■

Proof of Proposition 4. (i) Under $E(\bar\beta)$ the firm's shortfall is $(x-k)_+$ when $|\beta| \le \bar\beta$ and $(x+\beta-k)_+/(1+\gamma\phi(x))$ when $|\beta| > \bar\beta$, so its first-order condition is

$$c \;=\; \psi\,\mathbb{E}\Big[\mathbf{1}\{|\beta| \le \bar\beta\}\,(x-k)_+ \;+\; \mathbf{1}\{|\beta| > \bar\beta\}\,\frac{(x+\beta-k)_+}{(1+\gamma\phi(x))^2}\Big] \;=:\; MB_E(k;\bar\beta).$$

Because the trigger does not depend on k, the right-hand side is a continuous, strictly decreasing function of k for each $\bar\beta$ on the range where it is positive. As $\bar\beta \to \infty$ the second indicator vanishes almost surely and $MB_E \to MB_R$; as $\bar\beta \to 0$ the first vanishes and $MB_E \to MB_D$, the discretionary marginal benefit with the shift included in the supervisor's estimate. The limits of k_E follow. Raising $\bar\beta$ moves mass from the second integrand to the first; the first exceeds the second pointwise in x exactly when $(1+\gamma\phi(x))^2(x-k)_+ \ge \mathbb{E}_\beta[(x+\beta-k)_+ \mid |\beta| > \bar\beta]$, and the displayed condition implies this for every conditioning set. Under the condition MB_E is non-decreasing in $\bar\beta$ for each k, and monotone comparative statics give k_E non-decreasing with $k_D \le k_E \le k_R$. The condition fails only where $\gamma\phi(x)$ is small and x is close to k, where the shift raises the expected shortfall by convexity more than the accommodation discount lowers it; in the calibration of Section 3.7 the monotonicity holds throughout (Figure 3, Panel B).

(ii) Suppose invocation occurs on an event $A(k)$ whose probability is decreasing in k and that in the invoked states the standard is $s_D(x+\beta, k)$. The firm's marginal benefit of capital then includes the term $-\psi\,\partial\Pr(A)/\partial k$ times the reduction in shortfall cost obtained in the invoked states, which is negative, so MB is lower than under any exogenous trigger for every k, and in the limit in which the firm can trigger invocation whenever it is short, the first-order condition is that of discretion and $k = k_D$.

(iii) Write $L_E(\bar\beta)$ for expected supervisory loss under $E(\bar\beta)$, including the invocation cost. As $\bar\beta \to \infty$ the trigger set has vanishing measure, $k_E \to k_R$, and $L_E \to L_R$. As $\bar\beta \to 0$ the trigger set has full measure, $k_E \to k_D$, and $L_E \to L_D + \kappa$. For the claim about large finite thresholds, let $\mu(\bar\beta) = \Pr(|\beta| > \bar\beta)$ and consider the effect of a small reduction in $\bar\beta$ from a large value, which admits to the trigger set states with $|\beta|$ close to $\bar\beta$. The direct effect per unit of measure admitted is $-(\frac{1}{2}\bar\beta^2 - \kappa) + A(\bar\beta)$, where A is the accommodation loss incurred in the admitted states and is bounded by a quantity of order $\gamma\phi(x)(x - k_R)_+^2$, uniformly in $\bar\beta$. The indirect effect runs through k_E: the change in MB_E per unit of measure admitted is of order $\mathbb{E}[(x + \beta - k)_+]/(1 + \gamma\phi)^2 - (x - k)_+$, which grows at most linearly in $\bar\beta$, and the loss consequence of the induced change in k is bounded by the marginal shortfall loss, which is finite. The direct gain grows quadratically in $\bar\beta$ while the direct accommodation loss is bounded and the indirect loss grows at most linearly, so for $\bar\beta$ large enough the total effect of admitting the extreme states is negative: $dL_E/d\mu < 0$ at small μ, and $L_E(\bar\beta) < L_R$ for all sufficiently large finite $\bar\beta$. The comparison with L_D is not sign-determined in general: at the threshold just found L_E is close to L_R, so $L_E < L_D$ if and only if $L_R < L_D$ there, which holds when the accommodation loss exceeds $\tau^2/2$; where it does not, the ranking at the best threshold must be computed, as in Table 5. ■

Proof of Proposition 5. Firm i chooses $m_i \geq 0$ to minimize $\frac{\psi}{2}(s - k_i)_+^2 + \frac{a}{2}m_i^2$ where $s = \sum_j w_j(x_j - m_j)$, taking the other firms' choices as given. The first-order condition is $-\psi\, w_i\, (s - k_i)_+ + a\, m_i = 0$, so $m_i^* = \psi\, w_i\, (s - k_i)_+/a$. Substituting into s, the displacement of the standard is $\sum_i w_i m_i^* = (\psi/a)\sum_i w_i^2 (s - k_i)_+$. For symmetric firms $(s - k_i)_+ = (s - k)_+$ for all i, and writing $\bar x = \sum_i w_i x_i$ for the unmanipulated statistic, the fixed point for $s > k$ solves $s = \bar x - (\psi/a)H(s - k)$, giving

$$\bar x - s \;=\; \frac{(\psi/a)\, H}{1 + (\psi/a)\, H}\, (\bar x - k)_+,$$

which is increasing in H, vanishes as $H \to 0$, and attains $(\psi/a)(\bar x - k)_+/(1 + \psi/a)$ at $H = 1$. For asymmetric firms the first-order displacement is the weighted sum displayed in the proposition, and the fixed point retains the dependence on $\sum_i w_i^2$. ■

Proof of Proposition 6. On $[0, \delta^*]$ the second term of (6) is $\frac{1}{2}(1 - q)\bar\delta^2$, so $L'(\bar\delta) = -\mathbb{E}[(|u| - \bar\delta)_+] + (1 - q)\bar\delta$ and $L''(\bar\delta) = \Pr(|u| > \bar\delta) + (1 - q) > 0$; the loss is strictly convex, $L'(0) = -\mathbb{E}|u| < 0$, and $L'(\bar\delta) \to (1 - q)\bar\delta > 0$ as the tail vanishes, so a unique interior root exists and is the minimizer on the interval. Writing $G(\bar\delta, q) = \mathbb{E}[(|u| - \bar\delta)_+] - (1 - q)\bar\delta$, implicit differentiation gives $d\bar\delta^*/dq = \bar\delta^*/(\Pr(|u| > \bar\delta^*) + 1 - q) > 0$. Because $(|u| - \bar\delta)_+$ is convex in u, a mean-preserving spread raises $\mathbb{E}[(|u| - \bar\delta)_+]$ at every $\bar\delta$ and so raises the root. For $\bar\delta > \delta^*$ the capture term is constant at $\frac{1}{2}(1 - q)\delta^{*2}$ and the truncation term decreases to zero, so $\inf_{\bar\delta > \delta^*} L = \frac{1}{2}(1 - q)\delta^{*2}$, attained only in the limit of no bound; the interior solution is the global optimum if and only if $L(\bar\delta^*)$ lies below that value, and since $L(\bar\delta^*) \leq L(\delta^*) = \frac{1}{2}\mathbb{E}[(|u| - \delta^*)_+^2] + \frac{1}{2}(1 - q)\delta^{*2}$ the condition holds strictly whenever the truncation term at δ^* is positive, which it is whenever δ^* lies inside the support of $|u|$, and holds a fortiori as δ^* grows relative to the dispersion of u. ■

Appendix B: Calibration and Replication

Table 4. Central calibration.

Parameter	Symbol	Value	Role
Mean of the signal	$\mathbb{E}[x]$	1.00	normalization
Coefficient of variation of the signal	–	0.30	lognormal dispersion
Shortfall weight	γ	0.75	supervisor's loss (2), with $\phi(x) = x$
Late-capital cost	ψ	3.00	firm's cost of closing a shortfall
Unit cost of advance capital	c	0.30	firm's cost of k
Conditional s.d. of the state	σ_{ε}	0.20	enters all regimes as $\sigma_{\varepsilon}^2/2$
S.d. of the non-contractible shift	τ	0.15	information the mapping omits
Invocation cost	κ	0.010	cost of using the escape
Capture: distortion per period	b/a	1.00	Figure 2
Capture: reversal probabilities	q_D; q_R	0.10; 0.30, 0.50, 0.70	Figure 2
Capture: bound	$\bar{\delta}$	2.50	Figure 2
Manipulation: cost convexity	a	4.00	Figure 4

The signal is lognormal with the stated mean and coefficient of variation; the shift is normal with mean zero and standard deviation τ, independent of the signal. Expectations are computed by quadrature on a product grid of 1,201 signal points and 201 shift points. The firm's first-order condition is solved by bracketing. Under the escape regime the threshold is chosen on a grid of 61 points spanning $[0, 4.5\tau]$, and the reported optimum is the grid minimum. The code, in Python with NumPy, SciPy and Matplotlib, produces Figures 1 to 4 and Tables 4 and 5 from a single script and is supplied with the paper.

The empirical results of Section 4.5 use the sheet "Calculations" of EIOPA's workbook for end-September 2026, which gives for every calendar day from 30 June 1991 to 30 September 2026 the 36-month average of the EIOPA equity index, the proportional deviation of the current level from it, the unbounded adjustment and its truncation to ± 10 points, together with the daily index levels on trading days. Month-end values are the published monthly standards. The unbounded series is used as the distribution of u in Proposition 6; the stop-loss premium at a bound is its sample mean of $(|u| - \bar{\delta})_+$, the counterfactual charge is 39 points plus the truncation of u at the alternative bound, and the rationalizing q is $1 - \mathbb{E}[(|u| - \bar{\delta})_+]/\bar{\delta}$. Sampling uncertainty is assessed with a circular block bootstrap of the month-end series, with blocks of 36 months to match the averaging window that induces the series' dependence, 5,000 replications for the stop-loss premiums and 2,000 for the optimal widths, and 90 per cent intervals taken from the 5th and 95th percentiles of the replicates. A second script reads the extracted series and produces Figures 5 and 6 and Tables 6 and 7; the extracted series and the script are supplied with the paper.

Table 5. Sensitivity of the main results to the calibration.

Case	k_R	k_D	k_D/k_R	L_R	L_D	L_E^*	$\bar{\beta}^*$	invocation frequency
central	1.040	0.665	0.639	0.0607	0.0753	0.0574	0.203	0.177
$\gamma = 0.5$	1.040	0.780	0.750	0.0509	0.0510	0.0468	0.180	0.233
$\gamma = 1.0$	1.040	0.556	0.535	0.0705	0.1066	0.0677	0.203	0.177
$\tau = 0.05$	1.040	0.652	0.627	0.0507	0.0737	0.0507	0.184	0.000
$\tau = 0.25$	1.040	0.691	0.665	0.0807	0.0780	0.0635	0.225	0.368
$\kappa =$ 0.003	1.040	0.665	0.639	0.0607	0.0753	0.0558	0.169	0.251
$\kappa =$ 0.030	1.040	0.665	0.639	0.0607	0.0753	0.0596	0.281	0.062
CV $= 0.20$	0.958	0.679	0.709	0.0469	0.0583	0.0438	0.203	0.177
CV $= 0.40$	1.148	0.648	0.565	0.0838	0.1000	0.0801	0.203	0.177
$\psi = 2.0$	0.932	0.498	0.535	0.0753	0.1084	0.0729	0.236	0.115
$\psi = 4.0$	1.112	0.762	0.686	0.0535	0.0607	0.0497	0.191	0.200

Each row changes one parameter from the central calibration. L_E^* is the expected loss at the best threshold $\bar{\beta}^*$. In the row $\tau = 0.05$ the best threshold lies at the top of the grid and the escape is effectively never invoked, so $L_E^* = L_R$.

Declarations

No funding was received for this study. The author declares no competing interests. The empirical results use EIOPA's published workbook on the symmetric adjustment of the equity capital charge (end-September 2026), which is publicly available; the extracted series and the scripts that produce every figure and table are supplied with the paper.

References

Acharya, V. V., & Yorulmazer, T. (2007). Too many to fail: An analysis of time-inconsistency in bank closure policies. *Journal of Financial Intermediation*, 16(1), 1–31.

Actuarial Association of Europe. (2025). *Solvency II 2020 review* (The European Actuary No. 44). Actuarial Association of Europe.

Adrian, T., & Shin, H. S. (2010). Liquidity and leverage. *Journal of Financial Intermediation*, 19(3), 418–437.

Aegon Asset Management. (2022). *Symmetric adjustment of the equity capital charge under Solvency II.* Aegon Asset Management.

Aslan, I. (2026a). *Underwriting cycles and the calibration of premium risk under Solvency II.* Working paper, Zürich.

Aslan, I. (2026b). *Existential risk from artificial intelligence and the limits of insurability.* Working paper, Zürich, October 2026; prepared for the Workshop on the Measurement, Modeling, and Management of Existential Risk (from AI), Georgia State University, Atlanta, February 2027.

Aslan, I. (2026c). *Conventions, uncertainty, and financial instability: The reinsurance underwriting cycle through a post-Keynesian lens.* Working paper, Zürich.

Aslan, I. (2026d). Monetary sovereignty and catastrophe fiscal capacity: Turkey, Japan, and the limits of disaster self-insurance. *Review of Political Economy.* Advance online publication. https://doi.org/10.1080/09538259.2026.2724127

Aslan, I. (2026e). *Price softening, structural discipline: The reinsurance underwriting cycle after alternative capital.* Working paper, Zürich.

Barro, R. J., & Gordon, D. B. (1983). Rules, discretion and reputation in a model of monetary policy. *Journal of Monetary Economics*, 12(1), 101–121.

Bernanke, B. S., & Mishkin, F. S. (1997). Inflation targeting: A new framework for monetary policy? *Journal of Economic Perspectives*, 11(2), 97–116.

Boot, A. W. A., & Thakor, A. V. (1993). Self-interested bank regulation. *American Economic Review*, 83(2), 206–212.

Carpenter, D., & Moss, D. A. (Eds.). (2014). *Preventing regulatory capture: Special interest influence and how to limit it.* Cambridge University Press.

Cummins, J. D., & Outreville, J. F. (1987). An international analysis of underwriting cycles in property-liability insurance. *Journal of Risk and Insurance*, 54(2), 246–262.

Dal Bó, E. (2006). Regulatory capture: A review. *Oxford Review of Economic Policy*, 22(2), 203–225.

Danielsson, J., Embrechts, P., Goodhart, C., Keating, C., Muennich, F., Renault, O., & Shin, H. S. (2001). *An academic response to Basel II* (Special Paper No. 130). LSE Financial Markets Group.

Danielsson, J., James, K. R., Valenzuela, M., & Zer, I. (2016). Model risk of risk models. *Journal of Financial Stability*, 23, 79–91.

Davidson, P. (1991). Is probability theory relevant for uncertainty? A post Keynesian perspective. *Journal of Economic Perspectives*, 5(1), 129–143.

Dequech, D. (2011). Uncertainty: A typology and refinements of existing concepts. *Journal of Economic Issues*, 45(3), 621–640.

EIOPA. (2020). *Opinion on the 2020 review of Solvency II* (EIOPA-BoS-20/749). European Insurance and Occupational Pensions Authority.

EIOPA. (2021). *Artificial intelligence governance principles: Towards ethical and trustworthy artificial intelligence in the European insurance sector.* Report of the Consultative Expert Group on Digital Ethics in Insurance. European Insurance and Occupational Pensions Authority.

EIOPA. (2026). *Symmetric adjustment of the equity capital charge: End of September 2026* [Data set]. European Insurance and Occupational Pensions Authority, published 5 October 2026.

European Commission. (2021). *Proposal for a Directive of the European Parliament and of the Council amending Directive 2009/138/EC as regards proportionality, quality of supervision, reporting, long-term guarantee measures, macro-prudential tools, sustainability risks, group and cross-border supervision* (COM(2021) 581 final). European Commission.

European Systemic Risk Board. (2015). *Report on systemic risks in the EU insurance sector*. European Systemic Risk Board.

Geraats, P. M. (2002). Central bank transparency. *Economic Journal*, 112(483), F532–F565.

Grace, M. F., & Leverty, J. T. (2012). Property-liability insurer reserve error: Motive, manipulation, or mistake. *Journal of Risk and Insurance*, 79(2), 351–380.

Harrington, S. E., & Danzon, P. M. (1994). Price cutting in liability insurance markets. *Journal of Business*, 67(4), 511–538.

Insurance Sweden. (2021). *Insurance Sweden on EIOPA's proposal to change the corridor for the symmetric adjustment*. Svensk Försäkring, 19 March 2021.

Kane, E. J. (1989). *The S&L insurance mess: How did it happen?* Urban Institute Press.

Keynes, J. M. (1921). *A treatise on probability*. Macmillan.

Keynes, J. M. (1936). *The general theory of employment, interest and money*. Macmillan.

Keynes, J. M. (1937). The general theory of employment. *Quarterly Journal of Economics*, 51(2), 209–223.

Knight, F. H. (1921). *Risk, uncertainty and profit*. Houghton Mifflin.

Kydland, F. E., & Prescott, E. C. (1977). Rules rather than discretion: The inconsistency of optimal plans. *Journal of Political Economy*, 85(3), 473–492.

Laffont, J.-J., & Tirole, J. (1991). The politics of government decision-making: A theory of regulatory capture. *Quarterly Journal of Economics*, 106(4), 1089–1127.

Lamm-Tennant, J., & Weiss, M. A. (1997). International insurance cycles: Rational expectations/institutional intervention. *Journal of Risk and Insurance*, 64(3), 415–439.

Lewis, D. K. (1969). *Convention: A philosophical study*. Harvard University Press.

Lohmann, S. (1992). Optimal commitment in monetary policy: Credibility versus flexibility. *American Economic Review*, 82(1), 273–286.

Mailath, G. J., & Mester, L. J. (1994). A positive analysis of bank closure. *Journal of Financial Intermediation*, 3(3), 272–299.

Orléan, A. (1999). *Le pouvoir de la finance*. Odile Jacob.

Orléan, A. (2014). *The empire of value: A new foundation for economics* (M. B. DeBevoise, Trans.). MIT Press.

Petroni, K. R. (1992). Optimistic reporting in the property-casualty insurance industry. *Journal of Accounting and Economics*, 15(4), 485–508.

Plantin, G., Sapra, H., & Shin, H. S. (2008). Marking-to-market: Panacea or Pandora's box? *Journal of Accounting Research*, 46(2), 435–460.

Power, M. (1997). *The audit society: Rituals of verification*. Oxford University Press.

Richman, R. (2021). AI in actuarial science – a review of recent advances – part 1. *Annals of Actuarial Science*, 15(2), 207–229.

Rogoff, K. (1985). The optimal degree of commitment to an intermediate monetary target. *Quarterly Journal of Economics*, 100(4), 1169–1189.

Stigler, G. J. (1971). The theory of economic regulation. *Bell Journal of Economics and Management Science*, 2(1), 3–21.

Legal instruments. Directive 2009/138/EC of the European Parliament and of the Council of 25 November 2009 on the taking-up and pursuit of the business of Insurance and Reinsurance (Solvency II), Articles 101 and 106. Commission Delegated Regulation (EU) 2015/35 of 10 October 2014 supplementing Directive 2009/138/EC, Articles 39, 117 and 172 and Annex II. Commission Delegated Regulation (EU) 2019/981 of 8 March 2019 amending Delegated Regulation (EU) 2015/35. Directive (EU) 2025/2 of the European Parliament and of the Council of 27 November 2024 amending Directive 2009/138/EC as regards proportionality, quality of supervision, reporting, long-term guarantee measures, macro-prudential tools, sustainability risks and group and cross-border supervision, OJ L, 8 January 2025. Commission Delegated Regulation (EU) 2026/269 amending Delegated Regulation (EU) 2015/35, OJ L, 18 February 2026. Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act), Annex III and Articles 9 to 14.