International Journal of Computer Information Systems and Industrial Management Applications
ISSN 2150-7988 Volume 18(16s) 2026 pp. 391-404
https://doi.org/10.70917/ijcisim-2026-4589
Article

Encoding Statutory Variance in Property Rating Systems: A Mandate-Strength Classification and Specification Method for Multi-State Catastrophe Mitigation Credits

Anushka V. Rodi[1], Animesh Dhole[2], Harpreet Singh Siddhu[3], Sravan Reddy Kathi[4]

[1] Independent Researcher, Madison, Wisconsin, USA.
Email: anushkarodi@gmail.com
ORCID: 0009-0005-0652-1725
[2]Independent Researcher, Little Rock, Arkansas, USA.
Email: animeshudd@gmail.com
ORCID: 0009-0004-3688-7130
[3] Independent Researcher, Warrington, Pennsylvania, USA.
Email: hpsiddhu@gmail.com
ORCID: 0009-0008-5397-7495
[4] Independent Researcher, Bridgeport, Pennsylvania, USA.
Email: sravanreddykathi55@gmail.com
ORCID: 0009-0000-7646-3085

Abstract: Property insurance rating engines encode jurisdiction-specific catastrophe mitigation credits as local configuration, with no machine-resolvable link to the statutory provisions that require them. Practitioner discussion of this problem generally treats inter-state variation as a matter of differing credit percentages, to be handled by parameterising a single rule. We show that this framing is incorrect. Examining the operative catastrophe mitigation credit provisions of five U.S. jurisdictions from primary sources, we find that obligations differ not in magnitude but in kind: Florida imposes the credit in mandatory terms on the rate filing itself; New York's mandate is delegated, with the statute binding the regulator and the implementing regulation binding the carrier under trigger conditions absent from the statute; California's obligation binds conditionally, only on insurers whose property rating plans segment on wildfire risk, and is mandatory in content once that gate is passed; Texas authorises a construction-system discount without requiring it in the voluntary market, while its building-code mitigation credits bind only the residual-market association through a manual adopted by reference; and Ohio's property rate filing provisions contain no analogous mandate. Four obligation classes and two further structural axes -instrument layer and bound population -therefore appear across five jurisdictions on a single compliance dimension. From this finding we derive three contributions. First, a mandate-strength classification -mandatory, conditional, permissive, absent -stated as a triple over applicability, content strength, and bound population, so that obligation class becomes a first-class property of a rating rule rather than an implicit assumption. Second, a specification method encoding each obligation as a rule object that carries its enabling and operative instruments separately, together with version, mandate class, trigger, adjustment target set, and ordering constraint, and binds to core platform primitives including SAP BRFplus decision tables. Third, a compliance index whose satisfaction predicate is decidable from carrier-internal state and therefore does not presuppose the correctness oracle it exists to approximate, with its alert threshold derived from carrier loss tolerance rather than asserted. A worked analytical example over a synthetic 20,000-policy multi-state portfolio quantifies the consequence of ignoring mandate class: a uniform parameterised configuration diverges from the statute-faithful encoding on 11,875 policies (59.4%), and the divergence is driven by trigger mismatch rather than factor mismatch by a ratio of 22:1 in value. The statutory analysis is the paper's empirical core and is verifiable against the cited enacting texts; the portfolio is synthetic and the quantitative results are analytical rather than observed.

Keywords: Property insurance ratemaking, regulatory compliance, statutory encoding, catastrophe mitigation credits, rating logic specification, core insurance systems, decision tables, SAP BRFplus, rules as code, compliance verification, insurance regulation technology.

1. Introduction

A rating factor stored in a decision table is, to the platform executing it, a number. Nothing in the artifact records which statutory provision requires that number, which version of that provision it reflects, or whether the provision has since been amended. Modern policy administration platforms -Guidewire, Duck Creek Policy, and SAP core environments with the Business Rule Framework plus (BRFplus) [1] -provide expressive and well-tooled configuration surfaces, and none of them provides this linkage. The obligation is discharged by an analyst writing a requirement, a developer implementing it, and a release process deploying it; the connection between the legal instrument and the executing artifact survives only in project documentation, if anywhere.

The conventional response treats this as a documentation problem with a documentation solution. We argue it is a *modelling* problem, and that the modelling error is more specific than is generally recognised. When practitioners describe multi-state rating complexity, the description is almost always quantitative: different states require different credit percentages, so the rating engine holds a table of factors keyed by state. This framing is convenient because it maps cleanly onto existing configuration surfaces -a rate table with a state dimension is trivial to build.

It is also wrong. Section III examines the operative catastrophe mitigation credit provisions of five jurisdictions and finds that they differ in the *nature* of the obligation, not merely its size. One state requires the credit outright in the rate filing. A second requires it, but through a two-step delegation in which the statute binds the regulator and the implementing regulation binds the carrier under conditions the statute never states. A third requires it only if the insurer's own rating plan takes a particular form, making the carrier's configuration part of the trigger, and then prescribes a closed list of factors the plan must reflect. A fourth permits a discount without requiring it in the voluntary market, while imposing mitigation credits on its residual-market association alone. A fifth imposes no such obligation at all. A factor table keyed by state cannot represent any of this, because the thing that varies is not the factor.

This paper takes that finding as its foundation and builds a specification method on it. Section IV defines a mandate-strength classification. Section V gives the encoding schema. Section VI defines a compliance index that is computable without knowing in advance what the legally correct rate would have been -a constraint that rules out the obvious definition. Section VII quantifies, through a worked analytical example, what a uniform parameterised configuration costs relative to a statute-faithful one. Section VIII gives the execution-layer binding to core platform primitives.

A. Scope and claims

We are explicit about the evidentiary status of each part. The statutory analysis in Section III is primary-source research and is verifiable: every provision cited was checked against its enacting text, and a reader can confirm each cell of Table I independently. The classification and specification method of Sections IV and V are design contributions. The quantitative results in Section VII are analytical: they are computed from the encoded register against a synthetic portfolio, and they establish what uniform parameterisation costs given real regulatory variance. They are not measurements of any carrier's production configuration, and we make no claim about how prevalent that configuration pattern is in practice.

B. Contributions

- **A verified five-jurisdiction statutory analysis** (Section III) establishing that catastrophe mitigation credit obligations differ in kind rather than magnitude, with four obligation classes and two further structural axes -instrument layer and bound population -observed across five jurisdictions on one compliance dimension.
- **A mandate-strength classification** (Section IV) -mandatory, conditional, permissive, absent -stated as a triple over applicability, content strength, and bound population, with encoding consequences and directional failure modes derived for each class.
- **A specification schema and ordering discipline** (Section V) representing obligations as versioned, provenance-carrying rule objects that separate the enabling instrument from the operative instrument binding

the carrier, admit disjunctive adjustment targets, and impose an explicit precedence relation between mandated and discretionary adjustments.

- **An oracle-free compliance index** (Section VI) with a derived rather than asserted alert threshold.
- **A binding contract to core platform primitives** (Section VIII), including a decision-table representation for SAP BRFplus.

2. Background and Related Work

A. Rate regulation and the filing pipeline

U.S. property and casualty insurance is regulated at state level, and the mechanism by which rating logic becomes lawful differs by jurisdiction. In **prior-approval** regimes the commissioner must approve a rate before it may be implemented; California's Proposition 103 framework operates this way for homeowners and other property and casualty lines, with a rate application deemed approved sixty days after public notice subject to enumerated conditions and in any event 180 days after receipt [2]. In **file-and-use** regimes a filed rate may be used from its effective date; the Texas Department of Insurance describes Texas as a file-and-use state, with residential and commercial property rates governed by Insurance Code Chapter 2251 [3], while the implementing rule provides that filed information may be used from the filing date *unless* the insurer is subject to prior approval under Subchapter D [4].

This distinction matters for any proposal to accelerate rate logic deployment. In a prior-approval regime, deploying a rate change before approval is itself a violation; the interval before approval is regulator-controlled and is not a deployment inefficiency to be optimised away. The correct objective in such a regime is deployment *readiness* on the approved effective date -neither early nor late. Where an obligation arises from a legislative mandate with a fixed statutory effective date and no approval step, no such constraint applies and implementation latency is unambiguously the carrier's own. We return to this in Section IX.

Filing submission itself is largely handled through the System for Electronic Rate and Form Filing (SERFF), the NAIC platform through which carriers submit rate, rule, and form filings and through which regulators manage review and disposition [5]. SERFF establishes what was filed and what was approved, with what effective date. It has no visibility into whether the carrier's rating engine subsequently executes the approved logic. The gap between an approved filing and a verified executing rule is the territory this paper addresses, and Section IX.C discusses how the two connect.

B. Machine-readable legal knowledge representation

A mature standards literature addresses the representation of legal norms in machine-processable form, and we do not claim to originate this idea. Akoma Ntoso, an OASIS standard, provides an XML vocabulary for the structure and semantics of legislative, parliamentary, and judicial documents [6]. LegalRuleML extends rule markup with the deontic and defeasible constructs legal reasoning requires, and has been shown to support automated reasoning through translation into modal defeasible logic [7], [8]. The Object Management Group's SBVR and DMN specifications provide business-vocabulary and decision-table formalisms respectively, DMN with multiple certified conforming implementations [9], [10]. The Rules as Code programme, surveyed for the OECD by Mohun and Roberts, argues for governments publishing official machine-consumable versions of rules alongside their natural-language counterparts, and has been piloted by several governments [11]. Business process compliance research formalises regulatory constraints as reusable specification patterns [12], [13].

These standards are deliberately agnostic about execution backends, and that agnosticism is precisely what leaves the present problem open. None of them specifies how an encoded obligation binds to a rate routine, a rating worksheet, or a BRFplus decision table, and none supplies a domain-specific classification of obligation strength for insurance rating. Our contribution is narrower than the standards work and complementary to it: the register presented in Section III could be serialised in LegalRuleML without loss, and the classification and binding contributions would be unaffected. The relationship runs in both directions -Rules as Code addresses the supply side, where the rulemaker publishes machine-consumable rules, while this paper addresses the consumption side, where the regulated firm must encode obligations that have not been so published and must remain able to demonstrate that its executing configuration still corresponds to them.

C. Ratemaking and fairness research

A substantial computational literature addresses ratemaking accuracy and equity -discrimination-free pricing under causal frameworks [14], the nesting of classical actuarial structures into neural architectures [15]. This work is largely orthogonal to ours. It asks whether a rate is statistically defensible or equitable; we ask whether the rate the engine executes corresponds to the obligation currently in force. A rate can be well-calibrated, fair by a chosen statistical criterion, and non-compliant simultaneously, because it reflects a superseded or misapplied provision.

3. Statutory Analysis: Catastrophe Mitigation Credit Obligations

This section presents the paper's empirical core. For each of five jurisdictions we identify the operative provision governing catastrophe mitigation credits in personal lines property insurance, and characterise the obligation it creates. Every citation was verified against the enacting text.

A. Method and verification protocol

Jurisdictions were selected to span the range of catastrophe exposure and regulatory posture present in the U.S. personal lines property market: two high-exposure coastal wind states with active mitigation programmes (Florida, Texas), a wildfire-exposed prior-approval state (California), a mixed-exposure northeastern state (New York), and a low-catastrophe-exposure midwestern state as a control (Ohio).

For each jurisdiction the procedure was: (i) identify the codified chapter governing property rate filings; (ii) search that chapter, and the corresponding administrative code title, for provisions addressing mitigation, discounts, credits, rate differentials, or deductible reductions; (iii) retrieve the current codified text of every candidate provision from the official state publisher rather than from a secondary compilation; (iv) confirm that the provision governs property rather than automobile or another line; (v) confirm that the provision is **enacted and currently codified**, not proposed, superseded, or repealed; and (vi) record the operative verb and any attached conditions. Where a statute delegates to a regulator, the implementing regulation was retrieved and treated as a separate instrument. Negative findings were recorded only after steps (i)–(iii) returned no candidate, and the searched corpus is stated with the finding.

A note on method, and on why verification matters. This protocol was developed in response to two distinct failures in earlier drafts of this analysis, both instructive. In the first, four of five citations pointed to automobile insurance rating provisions; each cited section was real, correctly formatted, and plausibly named, and none governed property insurance. In the second, a single jurisdiction's cell paired a correctly formatted citation to an enacted section with operative language drawn from a differently numbered provision in a bill that was never enacted; the citation, the caption, and the quoted obligation each looked sound in isolation and belonged to three different instruments. Step (v) was added in consequence. We record both failures because they are small instances of the paper's own thesis: a citation attached to a rating rule decays silently, the failure mode does not announce itself, and detection requires a procedure rather than a reading.

B. Findings

Jurisdiction	Enabling instrument	Operative instrument binding the carrier	Obligation created	Bound population	Class
Florida	Fla. Stat. §627.0629(1) [16]	Same; implemented by Fla. Admin. Code R. 69O-170.017 [17]	A rate filing for residential property insurance MUST INCLUDE actuarially reasonable discounts, credits, or other rate differentials, OR appropriate reductions in deductibles, for properties on which qualifying windstorm mitigation fixtures or construction techniques have been installed or implemented.	All residential property insurers	Mandatory

California	Cal. Ins. Code §1861.05 [2]	Cal. Code Regs. tit. 10, §2644.9 [18]	§2644.9(a)(1): an insurer whose rating plan segments, creates a rate differential, or surcharges based on wildfire risk SHALL COMPLY. §2644.9(d)(1): no insurer shall use a rating plan that does not reflect an enumerated closed list of community-level and property-level mitigation factors. The gate is a fact about the insurer's own rating plan.	Insurers segmenting on wildfire risk	Conditional applicability, mandatory content
New York	N.Y. Ins. Law §2346(3), (4) [19]	11 NYCRR §160.6 [20]; see also DFS Circular Letter No. 3 (2024) [21]	The statute directs that the Superintendent SHALL PROVIDE for an actuarially appropriate premium reduction for property fitted or retrofitted with hurricane or storm shutters, and separately for hurricane-resistant laminated glass. The regulation binds the carrier: INSURERS SHALL PROVIDE the reduction where such measures are installed on ALL exterior wall openings and meet a 110 mph sustained wind or eight-foot storm surge standard.	All homeowners insurers and NYPIUA	Mandatory, delegated
Texas	Tex. Ins. Code ch. 2006, subch. A, §2006.002 [22]	Same. Separately, 28 Tex. Admin. Code §5.4501 adopts by reference the TWIA rate manual [23]	§2006.002: an insurer MAY grant a discount in homeowners premiums for property constructed with an insulating concrete form system, on written verification; the subchapter is captioned as an optional premium discount. Separately, the TWIA manual provides building-code compliance credits binding on the residual-market association only.	Voluntary market: elective. Residual market: TWIA only	Permissive (voluntary); Mandatory but scope-limited (residual)
Ohio	Ohio Rev. Code §§3935.03–3935.04 [24]	None identified	No jurisdiction-specific catastrophe mitigation credit mandate identified. Property rate filings are governed by general rating standards and filing procedure applicable to fire and marine insurers. Corpus searched: ORC ch. 3935 and Ohio Admin. Code Div. 3901 for the terms mitigation, discount, credit, rate differential, and deductible reduction.	—	Absent

TABLE I. Catastrophe mitigation credit obligations across five jurisdictions. Each citation was verified against the current codified enacting text per the protocol of Section III.A. Enabling and operative instruments are recorded separately because they diverge in three of the five jurisdictions.

C. Interpretation

Four obligation classes appear across five jurisdictions on a single compliance dimension, and two further structural axes appear alongside them. The differences are not differences of degree.

Florida and New York both mandate, but the mandate has different shape. Florida binds the rate filing directly and reaches windstorm mitigation fixtures and construction techniques generally. New York's statute binds the *Superintendent*, not the carrier; the carrier-binding obligation lives in 11 NYCRR §160.6, which adds trigger conditions the statute never states -the protective measures must be installed on all exterior wall openings and must meet a specified wind-speed or storm-surge standard. A register that cites only the statute encodes an obligation whose trigger it cannot see. This is the **instrument-layer** axis: the enabling instrument and the operative instrument binding the regulated firm are frequently not the same document, and only the latter carries the executable trigger.

California's obligation is the most structurally interesting, and it is not scalar. Section 2644.9(a)(1) does not bind all property insurers; it binds those whose rating plans use wildfire risk to differentiate, so the trigger for *applicability* is a fact about the carrier's own filed rating plan rather than about the insured risk. Once that gate is passed, however, subdivision (d)(1) is categorically mandatory and prescribes a closed enumerated list of community- and property-level mitigation factors the plan must reflect. The obligation is therefore conditional in applicability and mandatory in content. An encoding that treats it as a single scalar class will misapply it in one direction or the other, depending on the carrier's plan design.

Texas exhibits the third axis. In the voluntary market §2006.002 authorises rather than requires: an insurer *may* grant a discount for construction with an insulating concrete form system, and a Texas configuration that omits it entirely may be fully compliant. Conversely, a configuration that applies the discount where the carrier has not filed it is applying an unfiled rate differential -a defect of a different kind, and one that a compliance process oriented solely toward ensuring mandated credits *are* applied will not detect. Texas's substantive windstorm mitigation credits, by contrast, are carried in the TWIA rate manual adopted by reference into administrative code, and bind the residual-market association alone. The obligation is real, mandatory in form, and reaches no voluntary-market carrier. This is the **bound-population** axis, and it is orthogonal to strength: an obligation may be mandatory and still not reach the entity whose configuration is being audited.

Ohio provides the control case. Applying a mitigation credit in Ohio on the strength of a rule written for Florida is not a conservative error; it is a deviation from the filed rating plan.

The consequence for system design is direct: a rating configuration of the form “if the risk has qualifying mitigation, apply factor f(state)” cannot be correct across these five jurisdictions, because f is not the thing that varies. What varies is whether an obligation exists, which instrument creates it, what triggers it, whether the trigger is a property of the risk or of the carrier, and which population of carriers it reaches.

D. On credit magnitudes and adjustment targets

None of the mandatory provisions fixes a numerical credit. Florida requires that the differential be *actuarially reasonable*; New York requires that it be *actuarially appropriate*. The statutes constrain the existence, trigger, and basis of the differential; the magnitude derives from the carrier's filed and approved rating plan. Encoding a percentage as though it were statutory conflates two distinct sources of authority and produces a rule that is wrong whenever the filed plan changes. The schema in Section V keeps the two separable.

Florida additionally illustrates that the **adjustment target** may itself be disjunctive. Section 627.0629(1) is satisfied by discounts, credits, or other rate differentials *or* by appropriate reductions in deductibles -operations on different rating parameters with different downstream effects on premium and on loss transfer. A schema whose adjustment field names a single parameter cannot represent an obligation satisfiable by either of two structurally distinct operations, and will report a compliant configuration as defective whenever the carrier elects the alternative. The schema of Section V therefore carries an adjustment target set rather than a single target.

4. A Mandate-Strength Classification

We formalise the observation of Section III as a classification. For a compliance dimension δ and jurisdiction *j*, the obligation is characterised by a triple

$$m(\delta, j) = \langle a, c, \sigma \rangle \qquad (1)$$

where *a* is the applicability condition under which the obligation attaches (unconditional, or a guard over risk attributes or carrier configuration state), *c* is the content strength once applicable (mandatory, permissive, or absent), and *σ* is the bound population of regulated entities the obligation reaches. The four named classes below are the projections of this triple that occur in the five-jurisdiction register; the triple is retained because two of the five jurisdictions are not adequately described by a scalar class.

Class	**Definition**	**Encoding consequence**	**Failure mode if misclassified**
Mandatory	The instrument requires the differential, subject only to conditions on the insured risk. Applicability is unconditional over the bound population.	Rule must be present, must execute before discretionary adjustments, and its absence is a compliance defect.	Omission under-credits and exposes the carrier to refund liability.
Conditional	The instrument requires the differential only where a stated condition holds, and that condition may concern the carrier's own rating plan rather than the risk. Content strength once applicable is separately specified and is typically mandatory.	Trigger must include the conditioning fact as a guard term; the guard is evaluated against carrier configuration state, not policy data alone. Content obligations are encoded as a closed factor list beneath the guard.	Applied unconditionally, the rule credits risks the provision does not reach. Gated correctly but with an open factor list, the rule under-reflects the prescribed factors.
Permissive	The instrument authorises but does not require the differential.	Rule is gated on carrier election. Absence is compliant; presence without a filed basis is not.	Treated as mandatory, produces an unfiled rate differential.
Absent	No instrument creates the obligation in this jurisdiction.	No rule. Any credit applied must derive from the filed plan alone.	Applying a credit imported from another jurisdiction deviates from the filed rate.

TABLE II. Mandate-strength classification with encoding consequences.

Three properties of this classification are worth noting. First, it is *orthogonal to magnitude*: two jurisdictions may mandate the same credit percentage under different classes, and two jurisdictions in the same class may differ entirely in magnitude. Configuration surfaces that parameterise magnitude by state therefore capture the wrong axis. Second, the failure modes differ in *direction*. Misclassifying mandatory as absent under-credits; misclassifying permissive or absent as mandatory over-credits. A compliance process designed around a single direction -typically, ensuring mandated credits are applied -is structurally blind to the other. Third, the *bound population is independent of strength*. The Texas residual-market credits are mandatory in form and reach no voluntary-market carrier; a register that records strength without scope will raise a false defect against every voluntary-market configuration in the state.

5. Specification Method

A. Rule object schema

Each obligation is encoded as a rule object carrying its provenance and class alongside its executable content. The schema fields are:

Field	Content	Purpose
Identifier	Stable identifier for the encoded obligation.	Referenced by audit artifacts and by binding metadata.
Jurisdiction, dimension	The state and compliance dimension addressed.	Scopes the rule and organises the register.
Enabling instrument	Citation to the instrument that creates the legal authority, where distinct from the operative one.	Preserves the delegation chain so that an amendment upstream is detectable even when the operative text is unchanged.
Operative instrument	Citation to the provision that binds the regulated firm and states the executable trigger.	Makes the linkage machine-resolvable rather than documentary, and ensures the trigger is read from the instrument that actually carries it.
Version	Version of each cited instrument that this rule encodes.	Makes staleness a comparison rather than a judgement.
Mandate class	The triple ⟨applicability, content strength, bound population⟩ of Section IV.	Determines gating, ordering, scope filtering, and the applicable failure modes.
Trigger	Condition under which the rule applies, over risk attributes and, for conditional obligations, carrier configuration facts.	Encodes the jurisdictional trigger rather than assuming a common one.
Adjustment target set	One or more (rating parameter, operation) pairs, any of which satisfies the obligation, with the elected member recorded.	The executable content; represents disjunctive satisfaction such as Florida's rate differential or deductible reduction. Magnitude sourced from the filed plan.
Precedence	Position in the ordering relation (Section V.B).	Guarantees mandated adjustments precede discretionary ones.
Audit template	Structure of the derivation record emitted on application.	Produces per-decision evidence at execution time.

TABLE III. Rule object schema.

The separation of *enabling instrument* and *operative instrument* from *adjustment* is what allows the distinct sources of authority identified in Section III to be maintained independently: an amendment to either instrument increments the corresponding version and invalidates the rule until re-encoded, while a change to the filed rating plan updates the adjustment magnitude without touching the statutory linkage. In the three jurisdictions where the two instruments differ, collapsing them into a single field loses either the delegation chain or the executable trigger.

B. Ordering discipline

Rules are evaluated in a defined precedence order, with mandated adjustments assigned strictly higher precedence than discretionary ones. This is a correctness requirement, not a convention. Filed rating manuals routinely state that a statutorily mandated credit applies to the base rate before discretionary schedule rating, and where the composed operations do not commute, the resulting premiums differ materially.

The point is easily missed because the common case hides it. Two multiplicative credits commute, so their order is immaterial and the discipline appears vacuous. Introduce a cap, a floor, or a suppression -all of which occur in filed plans and in statutory rate constraints -and order becomes numerically load-bearing. Take a wind-peril premium of $1,200, a statutorily mandated credit of 15%, and a discretionary programme cap that limits the post-adjustment premium to $900. Applying the credit first gives $1,200 × 0.85 = $1,020, which the cap then reduces to $900. Applying the cap first gives \$900, which the credit then reduces to \$765. The two orders differ by \$135, or 15.0% of the compliant premium, and only the first order satisfies the requirement that the mandated credit apply to the base rate. Assigning precedence by mandate class makes the correct order structural rather than a matter of developer discipline, and permits the ordering property to be checked before deployment rather than inferred from anomalous premiums afterwards.

C. Register consistency

Two rules are in conflict when they may apply simultaneously and write the same rating parameter. Restricting trigger conditions to a tractable fragment -membership of a categorical attribute in a finite set, and containment of an ordered attribute in a bounded interval, combined conjunctively and disjunctively -makes joint applicability decidable, so conflicts can be detected by construction over the register rather than discovered in production. The five-jurisdiction register presented here is conflict-free, as expected, since each rule is gated on its own jurisdiction and, in the Texas case, additionally on bound population; the check becomes substantive as dimensions are added within a jurisdiction.

6. Compliance Measurement

A. Why the obvious definition fails

The natural definition of a compliance index -agreement between executed premiums and legally correct premiums -is not computable. Determining the legally correct premium is the problem the specification method exists to address; an index defined in those terms presupposes its own solution. Any usable index must therefore be defined over properties decidable from state the carrier actually holds.

B. Definition

Let O_j denote the set of registered obligations applicable to jurisdiction j whose bound population includes the carrier. For each obligation o, define the satisfaction predicate $s(o) \in \{0, 1\}$ as the conjunction of three decidable conditions:

- **Bound** -a rule object encodes o and is bound to an execution primitive on the target platform.
- **Current** -the encoded version equals the operative version of each cited instrument in the register.
- **Verified** -the binding passed its most recent execution verification test.

The conjunction is deliberate: an obligation encoded but stale, or current but whose binding fails its test, yields no compliance value. Per-jurisdiction and aggregate scores follow:

$$R_j = (1 / |O_j|) \cdot \Sigma_{o+oj}\, s(o) \qquad (2)$$

$$w_j = P_j / \Sigma_k P_k \qquad (3)$$

$$\mathrm{RLVS} = \Sigma_j\, w_j \cdot R_j \qquad (4)$$

where P_j is written premium in jurisdiction j and the sum in (2) runs over all $o \in O_j$. Exposure weighting is material rather than cosmetic. Under a scenario in which a single Florida obligation becomes stale following an amendment, the exposure-weighted index reads 0.8975 against an unweighted mean of 0.9500 -a gap of 0.0525, because the degraded jurisdiction carries disproportionate premium. Under a compound scenario adding an unbound Texas obligation, the gap widens to 0.0962. An unweighted index systematically understates financial exposure, which is the quantity the index exists to track. The index also localises failure: in the single-stale scenario, Florida reports three of four obligations satisfied while the remaining jurisdictions report full satisfaction, and the failure reason is recoverable per obligation.

Propagation latency is deliberately excluded. Deployment speed is a process characteristic measured over an interval, while the index is a point-in-time statement about current configuration. Folding a retrospective process metric into a state index makes it uninterpretable: a carrier could improve its score through past deployment speed while currently executing stale logic. Deployment latency should be reported separately.

C. Threshold derivation

An alert threshold asserted without derivation is arbitrary. Let κ denote the mean relative premium impact of a single unsatisfied obligation, expressed as a fraction of premium. To first order, expected annualised exposure as a fraction of written premium is $(1 - \mathrm{RLVS}) \cdot \kappa$. A carrier setting tolerance τ on this quantity requires:

$$\theta = 1 - \tau / \kappa \qquad (5)$$

Tolerance τ	κ = 3%	κ = 5%	κ = 8%	κ = 10%
0.25%	0.917	0.950	0.969	0.975
0.50%	0.833	0.900	0.938	0.950
0.75%	0.750	0.850	0.906	0.925
1.00%	0.667	0.800	0.875	0.900

TABLE IV. Alert threshold θ as a function of carrier tolerance τ and mean per-obligation impact κ.

This reframes the threshold as a governance parameter the carrier sets deliberately, with a stated relationship to its own loss tolerance, rather than a constant inherited from a vendor default.

7. Worked Analytical Example

Method and status. To quantify the consequence of ignoring mandate class, we encode the five obligations of Table I under the schema of Section V and compare the resulting premiums against those produced by a uniform parameterised configuration. The portfolio is **synthetic** -20,000 policies across the five jurisdictions, generated under a fixed seed, with no real policyholder data at any stage. Mitigation attributes are assigned per jurisdiction to match the trigger actually specified by the operative instrument: windstorm mitigation fixtures and construction techniques for Florida, the enumerated wildfire factor list for California, shutters or laminated glass on all exterior wall openings for New York, and insulating concrete form construction for Texas. Credit magnitudes are illustrative filed values, consistent with Section III.D. The results below are therefore **analytical**: they establish what uniform parameterisation costs given the regulatory variance documented in Section III, and are not measurements of any carrier's production configuration.

The comparison configuration applies a single credit wherever any qualifying mitigation feature is present, at a common factor, in every jurisdiction -the pattern that follows naturally from treating inter-state variation as purely quantitative.

Jurisdiction	Policies	Statute-faithful ($)	Uniform config ($)	Difference ($)	Difference (%)
Florida	4,874	8,461,285	7,944,621	−516,664	−6.11
Texas	5,954	6,830,103	6,165,863	−664,241	−9.73
California	3,989	2,413,896	2,238,248	−175,648	−7.28
New York	2,874	2,183,110	1,995,955	−187,156	−8.57
Ohio	2,309	744,860	671,515	−73,344	−9.85
Total	**20,000**	**20,633,254**	**19,016,202**	**−1,617,052**	**−7.84**

TABLE V. Uniform parameterised configuration compared with the statute-faithful encoding. Wind-peril premium component only; not total homeowners premium.

The configurations diverge in **all five jurisdictions**, affecting **11,875 of 20,000 policies (59.4%)**. The mechanism differs by mandate class, exactly as Table II predicts. In Ohio, where no obligation exists, and in Texas, where the voluntary-market discount is permissive and unelected, the uniform rule applies a credit with no filed basis; in the Texas case it also silently imports a residual-market credit structure that does not bind the carrier at all. In Florida, California, and New York the credit is applied on the wrong trigger, reaching risks the operative provision does not.

A. Decomposition by error type

Separating divergence attributable to the *trigger* from divergence attributable to the *magnitude* isolates which modelling axis matters.

Error type	Policies affected	Value ($)	Share of total value
Trigger (credit applied under one configuration and not the other)	10,247	−1,547,899	95.7%
Magnitude (both configurations credit; filed factors differ)	1,628	−69,154	4.3%
Total	**11,875**	**−1,617,052**	**100%**

TABLE VI. Divergence decomposed by error type.

Trigger error exceeds magnitude error by a factor of **6.3 to 1 by policy count and 22 to 1 by value**. This is the substantive result of the example, and it corroborates the classification argument directly: the damaging modelling error is not applying the wrong percentage but applying the credit to the wrong population. A configuration surface that offers per-state factor parameterisation addresses 4.3% of the divergence by value and leaves the remainder untouched.

B. Error direction

Narrowing the comparison configuration -keying the credit on a single mitigation feature rather than any -produces divergence in both directions simultaneously: 4,588 policies credited where the operative provision does not require it, and 1,102 policies not credited where it does. Net divergence is $548,401 while gross divergence is $677,784, so the net figure understates by 19% as opposing errors partially cancel in aggregate.

This has a practical implication. A carrier reconciling realised premium against filed rates at portfolio level sees the smaller number, while every affected policy is independently mispriced and independently subject to challenge. Per-decision derivation records -the audit template of Table III -are the artifact that surfaces this; portfolio-level reconciliation is structurally incapable of it.

C. Sensitivity

Repeating the comparison across ten portfolio realisations of 20,000 policies each -the realisation reported in Table V together with nine further independently seeded portfolios -gives a mean relative divergence of −7.92% with a standard deviation of 0.06 percentage points, over a range of −8.05% to −7.84%. The divergence is therefore driven by the trigger structure encoded in the register rather than by portfolio composition, which is what the trigger-dominance result of Table VI would predict.

8. Execution-Layer Binding

The schema of Section V is platform-neutral. Statutory fidelity is achieved only when each rule object is bound to a concrete execution primitive in the carrier's core landscape, and the binding must preserve the fields that carry compliance meaning -enabling and operative instrument, version, mandate class, and precedence -rather than only the executable content.

Schema field	SAP BRFplus	Guidewire	Duck Creek Policy
Identifier, jurisdiction	Decision table name and version metadata	Rate routine identifier	Rating layer component ID
Enabling and operative instrument, version	Table documentation attributes retained for audit lineage, one per instrument	Routine metadata	Component metadata
Mandate class (applicability, strength, scope)	Guard condition column for applicability; ruleset membership for scope filtering	Routine eligibility conditions and product/state availability	Rating step conditions and line-of-business scoping

Trigger	Condition columns bound to context data objects	Routine eligibility conditions	Rating step conditions
Adjustment target set	Result columns carrying target parameter, operator, and value, one per admissible target, with the elected target flagged	Rate table factor, or deductible modification step	Rating layer output, or deductible step
Precedence	Ruleset priority and rule sequence position	Routine execution order	Step sequence
Audit template	Follow-up action emitting the derivation record	Routine audit output	Component trace output

TABLE VII. Binding contract across three core platforms. The BRFplus column maps to primitives documented in [1]. The Duck Creek columns map to the standard published concepts of each platform's rating architecture; the exact metadata carriers should be confirmed against the licensed documentation release in use before implementation.

Because the intermediate representation is uniform and only the target-binding metadata varies by platform, a single specification change propagates to heterogeneous backends without restatement. In an SAP core environment each rule object materialises as a versioned BRFplus decision table invoked from the rating calculation sequence: the trigger binds to condition columns evaluated against context data objects, the adjustment target set binds to result columns, the precedence binds to ruleset priority sequencing, and the audit template binds to a follow-up action emitting the derivation record to the compliance log at execution time. Both instrument citations are retained as table documentation attributes so that lineage survives export and re-import.

Status. The binding contract is specified and a decision-table representation is generated from the register for all encoded rules. Validation against a live rules engine -confirming import fidelity, priority handling, and follow-up action behaviour -has not been performed and is not claimed here. In particular, the correspondence between the precedence semantics of Section V.B and a given engine's native conflict resolution must be established per platform rather than assumed.

9. Discussion

A. The modelling axis

The central claim is that multi-state rating variance has been modelled on the wrong axis. Configuration surfaces optimise for parameterising magnitude by jurisdiction, because that is what a rate table with a state dimension naturally expresses. Section III shows that magnitude is not what principally varies, and Section VII quantifies the cost of the mismatch at 95.7% of divergence by value. Making mandate class, instrument layer, and trigger first-class fields addresses the dominant term.

B. Implications for rating logic testing

The direction asymmetry of Section VII.B has a consequence for test design that we consider the most immediately actionable finding. Rating test suites conventionally verify that a credit of the correct size is applied to a policy entitled to it. They rarely verify that the credit is *not* applied to a policy not entitled to it, because constructing that negative test requires knowing the jurisdictional trigger -precisely the knowledge an untraceable configuration fails to record. Encoding the trigger as a first-class field makes the negative test constructible, and the classification of Table II indicates which negative tests are required for each obligation class.

The classification yields a test obligation matrix directly. A mandatory obligation requires a positive test on the trigger population and a negative test on its complement. A conditional obligation requires, in addition, a pair of tests over the applicability guard: the rule must fire under a carrier configuration that opens the gate and must not fire under one that does not -a test that cannot be written at all unless carrier configuration state is representable in the fixture. A permissive obligation requires a negative test under non-election, which is the case most commonly omitted in practice. An absent obligation requires only a negative test, and a scope-limited obligation requires a negative test against carriers outside the bound population. This matrix is mechanically derivable from the register, which makes coverage of the negative direction auditable rather than dependent on test author judgement.

C. Relationship to filing infrastructure

The method is complementary to SERFF rather than competitive with it. SERFF establishes what was filed and approved and with what effective date; the register and index establish whether the approved logic is executing and remains current. A carrier operating both closes the loop: filing disposition drives version increments in the register, and verified binding state can be reconciled against filing records. Given the regime distinction of Section II.A, this coupling is also what allows deployment to be staged for readiness on an approved effective date in prior-approval jurisdictions, rather than accelerated indiscriminately. Comparable multi-condition, exception-bearing regulatory structures have been formalised as compliance patterns in adjacent regulated domains [12], [13], which suggests the classification is not specific to insurance rating; establishing that generalisation empirically is outside the scope of this paper.

10. Limitations

The statutory analysis of Section III is primary-source research and is independently verifiable against the cited enacting texts; the mandate-strength classification of Section IV follows from it directly. The limitations below therefore concern the quantification of Section VII and the generalisation of the method, not the statutory findings themselves.

Modelled quantification. The results of Section VII are computed from the encoded register against a synthetic portfolio and a constructed comparison configuration. They establish what uniform parameterisation costs given the documented statutory variance; they are not measurements of any carrier's production configuration, and we make no claim about how prevalent that pattern is in practice. Portfolio attribute distributions are plausible but uncalibrated to published market data, and credit magnitudes are illustrative filed values; the sensitivity analysis shows insensitivity to portfolio realisation but not to systematically different mitigation-adoption or geographic profiles. Divergence is expressed against the wind-peril premium component rather than total homeowners premium. The magnitude-error component of Table VI depends on the illustrative factors; the trigger-error component, which dominates by more than twenty to one in value, does not. Calibration to public sources -state regulator mitigation statistics and NAIC market share data -would strengthen the quantification and is the most valuable extension of this work.

Bounded statutory scope. The Ohio result is a negative finding bounded by the corpus stated in Table I: it establishes the absence of a mandate in the property rate-filing chapter and the corresponding administrative division, not the absence of any mitigation-related provision anywhere in Ohio law. The Texas permissive exemplar, §2006.002, addresses a specific construction system rather than a general mitigation certification programme; the class assignment is sound and the encoding consequence unchanged, but a jurisdiction with a broad permissive mitigation programme would be a stronger exemplar of the class.

Classification scope. The triple of Section IV was derived from and tested against one compliance dimension in five jurisdictions. Whether applicability, content strength, and bound population are sufficient axes for variable-restriction, disclosure, or notice obligations is untested. Extending the register requires per jurisdiction and dimension the same verification discipline applied in Section III, which we estimate at several hours of qualified analyst time per pair.

Platform validation. As stated in Section VIII, the binding contract has not been validated against a live production rules engine, and the Guidewire and Duck Creek columns of Table VII are derived from published architectural concepts rather than from implementation.

Open-textured standards. Obligations turning on terms such as *actuarially reasonable* are not fully codifiable. The method encodes their structural component -existence, trigger, precedence, provenance -while the substantive actuarial judgement remains with the actuary and outside the register.

11. Conclusion

Multi-state property rating compliance is conventionally modelled as a parameterisation problem: different states, different credit percentages, one rule. Examining the operative catastrophe mitigation credit provisions of five jurisdictions from primary sources, we find this framing does not survive contact with the statutes. Obligations differ in kind -mandatory in Florida on the rate filing itself, mandatory but delegated in New York with the executable trigger residing in the implementing regulation, conditional in California on a fact about the carrier's own rating plan and mandatory in content once that gate is passed, permissive in the Texas voluntary market while binding the residual market alone, and absent in Ohio -producing four obligation classes and two further structural axes on a single compliance dimension.

We derived from this a mandate-strength classification stated as a triple over applicability, content strength, and bound population; a specification schema separating enabling from operative instrument and admitting disjunctive adjustment targets; an ordering discipline for mandated versus discretionary adjustments; a compliance index computable without a correctness oracle and with a derived alert threshold; a binding contract to core platform primitives. A worked analytical example quantifies the cost of the conventional framing: a uniform parameterised configuration diverges on 59.4% of policies, with trigger error exceeding magnitude error by 22 to 1 in value, and with divergence running in both directions once the comparison trigger is narrowed.

The claims are deliberately bounded. The statutory analysis is primary-source research and is independently verifiable. The quantification is analytical and rests on a synthetic portfolio and constructed comparison configurations. Platform validation is future work. What the paper establishes is that the variance is real, that it is structural rather than numerical, and that a specification method addressing the structural axis is both constructible and measurable.

References

1. C. Ziegler and T. Albrecht, BRFplus -Business Rule Management for ABAP Applications. Boston, MA, USA: SAP PRESS, 2011.
2. Cal. Ins. Code §1861.05 (Proposition 103, Insurance Rate Reduction and Reform Act, 1988).
3. Tex. Ins. Code ch. 2251 (Rates). Austin, TX, USA: Texas Legislature.
4. 28 Tex. Admin. Code §5.9334 (Requirements for rate and rule filing submissions); Tex. Ins. Code ch. 2251, subch. D.
5. National Association of Insurance Commissioners, System for Electronic Rate and Form Filing (SERFF). Kansas City, MO, USA: NAIC.
6. OASIS LegalDocumentML (LegalDocML) Technical Committee, “Akoma Ntoso version 1.0, part 1: XML vocabulary,” OASIS Standard, 2018.
7. M. Palmirani, G. Governatori, T. Athan, H. Boley, A. Paschke, and A. Wyner, “LegalRuleML core specification,” OASIS LegalRuleML Technical Committee, OASIS Standard.
8. H.-P. Lam and M. Hashmi, “Enabling reasoning with LegalRuleML,” Theory and Practice of Logic Programming, vol. 19, no. 1, pp. 1–26, 2019.
9. Object Management Group, “Semantics of Business Vocabulary and Business Rules (SBVR),” OMG Available Specification, Needham, MA, USA.
10. Object Management Group, “Decision Model and Notation (DMN),” OMG Available Specification, Needham, MA, USA.
11. J. Mohun and A. Roberts, “Cracking the code: Rulemaking for humans and machines,” OECD Working Papers on Public Governance, No. 42, OECD Publishing, Paris, France, 2020. doi: 10.1787/3afe6ba5-en.
12. A. Elgammal, O. Turetken, W. J. van den Heuvel, and M. Papazoglou, “Formalizing and applying compliance patterns for business process compliance,” Software and Systems Modeling, vol. 15, no. 1, pp. 119–146, 2016.
13. G. Governatori and S. Sadiq, “The journey to business process compliance,” in Handbook of Research on Business Process Modeling. Hershey, PA, USA: IGI Global, 2009, pp. 426–454.
14. M. Lindholm, R. Richman, A. Tsanakas, and M. V. Wüthrich, “Discrimination-free insurance pricing,” ASTIN Bulletin, vol. 52, no. 1, pp. 55–89, 2022.
15. R. Schelldorfer and M. V. Wüthrich, “Nesting classical actuarial models into neural networks,” SSRN Electronic Journal, ID 3320525, 2019.
16. Fla. Stat. §627.0629 (Residential property insurance; rate filings).
17. Fla. Admin. Code R. 69O-170.017 (Windstorm mitigation discounts).
18. Cal. Code Regs. tit. 10, §2644.9 (Consideration of mitigation factors; wildfire risk models), eff. Oct. 14, 2022.
19. N.Y. Ins. Law §2346 (Reduction in rates of fire insurance, homeowners insurance or property/casualty premiums for residential property); see also §2346-a.
20. 11 NYCRR §160.6 (Premium reductions for loss prevention measures).
21. New York State Department of Financial Services, Insurance Circular Letter No. 3 (2024), Albany, NY, USA.
22. Tex. Ins. Code ch. 2006, subch. A, §2006.002 (Optional premium discount for use of insulating concrete form system).
23. 28 Tex. Admin. Code §5.4501 (Adoption by reference of the rate manual governing the writing of windstorm and hail insurance by the Texas Windstorm Insurance Association).
24. Ohio Rev. Code §§3935.03–3935.04 (Rating regulations; filing of rates and schedules).