

# Two Failure Modes Require Architectural Change: A Formal Harmonization Gap Analysis of the EU AI Act, Vietnam’s AI Law, and the ASEAN AI Governance Guide<sup>1</sup>

Ngo Thai Bao

With  
Apart Research

## Abstract

There is a common assumption in tech policy that the EU sets the global baseline for AI regulation, with other regions simply converging toward it. However, a formal analysis of Vietnam’s new AI Law (Law No. 134/2025/QH15)—Southeast Asia’s first binding national AI legislation, taking effect on March 1, 2026—shows this isn’t always true. In fact, for minimal-risk AI systems, Vietnam enforces a universal human-control principle where the EU AI Act remains completely silent, making Vietnam the stricter regime in terms of formal legal force.

I arrived at this finding by translating the EU AI Act, Vietnam’s AI Law, and the ASEAN AI Governance Guide into HSDL, a computable policy language. This allowed me to transform traditional, manual comparative law into a brute-force computable analysis across 2,880 distinct scenarios. In doing so, I identified four structural ways these frameworks clash: regulatory voids, bindingness gaps, partition incompatibility, and obligor mismatches. Across the seven pairs of obligations I tested, two of these failure modes (partition incompatibility and obligor mismatches) cannot be fixed by simply tweaking legal text; they require fundamental architectural redesigns.

With the ASEAN Digital Economy Framework Agreement nearing signature, and countries like Indonesia and Malaysia drafting their own rules, Vietnam has proven that emerging economies do not have to default to EU precedents. These findings carry direct implications for how Southeast Asia should structure its transition to binding AI governance.

## 1. Introduction

Typically, researchers compare AI governance frameworks by reading statutes and writing prose. While useful, this method doesn’t scale well and struggles to answer a very practical question: for a specific AI system in a specific scenario, exactly which rules differ between two countries, and by how much?

Vietnam’s AI Law and its implementing decree (Decree No. 142/2026/NĐ-CP, which rolls out staggered compliance deadlines through 2027) change the regional landscape and render previous comparative studies outdated. Vietnam is also not acting in a vacuum. Throughout 2026, Indonesia and

---

<sup>1</sup>Research conducted at the The Global South AI Safety Hackathon, June 2026

Malaysia have been advancing their own mandatory AI requirements, and the ASEAN Digital Economy Framework Agreement (DEFA) is expected to be signed this November. As the first ASEAN member to adopt a binding national AI law, Vietnam currently serves as the real-time reference point for its neighbors.

**The principal finding.** When I formally modeled these regulations, I found that Vietnam’s law is actually stricter than the EU AI Act for an entire risk tier—inverting the standard narrative of EU regulatory dominance. Specifically, at the “minimal risk” level, Vietnam’s Article 4(2) mandates human control universally. In contrast, the EU AI Act’s equivalent rule (Art. 14) is tier-gated and remains silent at this lowest risk level. I verified this divergence across 2,880 distinct configuration contexts.

By encoding all three frameworks into HSDL policies, I surfaced four fundamentally different ways harmonization fails:

1. **Regulatory voids:** One regime regulates an issue while the other is completely silent.
2. **Bindingness gaps:** Both regimes address the issue, but with different levels of legal force.
3. **Partition incompatibility:** The laws categorize the regulatory space in fundamentally incompatible ways.
4. **Obligor mismatch:** The laws agree on the rule’s strictness, but disagree on *who* is responsible for complying.

**My main contributions:** First, I extended my custom policy language, HSDL (Holy Seed Description Language), to handle legal harmonization by adding a bindingness-labeling layer, making cross-jurisdiction comparisons decidable. Second, I conducted the first formal three-way comparison of EU, Vietnamese, and ASEAN AI laws. Third, I introduced a new analytical dimension to measure “obligor mismatches,” showing that even when laws appear harmonized, they might target completely different actors. Finally, I ran a sensitivity analysis to separate findings that are robust across interpretations from those that depend heavily on how I encoded the law.

## 2. Related Work

Most comparative AI governance literature focuses on the US–EU–China triangle. When Southeast Asia is mentioned, it is usually framed as a region relying on voluntary guidelines, lagging behind the binding-law frontier. Lu and Tie (2025) provided the closest prior work with their ASEAN–EU comparison, but they relied on prose analysis and published before Vietnam enacted its binding law. Because Vietnam’s legislation is so recent, my three-way formal comparison—and the discovery that Vietnam is sometimes the stricter regime—has no precedent in the current literature.

There is a growing body of work that formalizes the EU AI Act itself. Teams like Guldemann et al. (COMPL-AI) and Hernandez et al. have successfully mapped EU requirements to benchmarks and ISO standards. However, those studies evaluate systems *within* the EU framework. My approach is comparative: I map multiple jurisdictions onto a single shared schema, allowing me to compute compliance gaps across borders.

Technically, my framework builds on the LegalRuleML and defeasible logic traditions (Governatori and Rotolo, 2010; Palmirani et al., 2011). I opted to use HSDL, a lighter, purpose-built language, trading the broad generality of LegalRuleML for strict decidability over this specific three-jurisdiction model. Additionally, I treat legal fragmentation as a computable “arbitrage” opportunity. While

Marwala (2026) discussed AI governance arbitrage conceptually, my method actually computes these differentials on a per-rule, per-configuration basis.

### 3. Formal Framework (HSDL)

I originally developed the Holy Seed Description Language (HSDL) to manage dynamic AI-agent authorizations. For this paper, I adapted it to model legal harmonization. (*Note: Readers primarily interested in policy takeaways can skip directly to Section 6*).

**Context space.** I define  $Ctx$  as a finite set of typed dimensions representing a specific AI scenario:

$$\begin{aligned} Ctx = & risk\_tier \times sector \times system\_role \times lifecycle\_stage \\ & \times interacts\_with\_human \times serious\_harm\_discovered \\ & \times modification\_increases\_risk \times existing\_sector\_certification \end{aligned}$$

This gives  $|Ctx| = 2,880$  unique points. A policy  $P = \{R_1, \dots, R_n\}$  is a set of named rules. By Lemma 1,  $S(R) = S_{num}(R) \times S_{cat}(R)$ , and  $Satisfying(P) = \bigcup_i S(R_i)$  (Theorem 2).

**Bindingness (H1–H6).** I measure a regime’s regulatory force using the output scale  $B = \{\perp, Voluntary, Recommended, Binding\}$ , ranked 0 to 3. A regime is defined as  $(P, \lambda)$  with  $\lambda : P \rightarrow B$ . Applying the *lex strictior* (strictest rule applies) principle:

$$\beta_{(P, \lambda)}(ctx) = \max\{\lambda(R_i) : R_i \in P, ctx \in S(R_i)\}$$

I define a **Regulatory Void** as  $Void \downarrow (A \rightarrow B) = \{ctx : \beta_A(ctx) > \perp \wedge \beta_B(ctx) = \perp\}$  — meaning regime A imposes a rule while regime B is entirely silent. A **Gap Set** is  $Gap_{A,B}(a, b) = \{ctx : \beta_A(ctx) = a \wedge \beta_B(ctx) = b\}$ , and the **Directional Shortfall** is  $Gap \downarrow (A \rightarrow B) = \bigcup_{a > b} Gap_{A,B}(a, b)$ . As shown in Corollary H5.1, this shortfall represents the exact space where a company minimizing compliance costs would prefer regime B over A.

**Classification schemes and partitions (H7–H8).** Beyond mere strictness, regimes can clash architecturally. H7 examines whether one regime’s risk categories can even be translated into another’s. H8 checks if the regions of  $Ctx$  covered by each regime cleanly map onto each other (Theorem C) — a structural issue completely independent of how strict the rules are.

**Obligor as an independent dimension (H9–H13).** Finally, let  $O$  be a finite, flat set of obligated actors:  $\{Provider, Deployer, Regulator, ThirdPartyCertifier\}$ . Each rule designates an obligor  $\omega : P \rightarrow \mathcal{P}(O)$ . When rules overlap, they aggregate by union. An **Obligor Gap** occurs when regimes agree on how strict a rule is, but disagree on *who* has to follow it:


$$ObligorGap(A, B) = \{ctx : \beta_A(ctx) = \beta_B(ctx) > \perp \wedge \omega_A(ctx) \neq \omega_B(ctx)\}$$

### 4. Policy Encodings: Seven Obligation Pairs

I evaluated seven specific areas of AI regulation. **Table 1** breaks down the four failure modes I found, and **Table 2** details the encodings for the seven obligation pairs.

**Table 1: Four Failure Modes of Harmonization**

| Failure mode              | Canonical example                                                            | Fixable by text amendment?                    |
|---------------------------|------------------------------------------------------------------------------|-----------------------------------------------|
| Regulatory Void           | G1: VN silent on post-market High-risk supervision — closed by Article 10(5) | Yes                                           |
| Bindingness Gap           | G3: VN binds at <code>risk_tier=Minimal</code> where EU is silent            | Yes — adjust either label                     |
| Partition Incompatibility | EU’s tier-only cover vs. VN’s multi-axis cover — neither refines the other   | <b>No</b> — needs re-architecting             |
| Obligor Mismatch          | 4 instances (G1, G3, G4, G6), 3 distinct mismatch types                      | <b>No</b> — needs obligor-allocation decision |



**Figure 1.** The four failure modes, grouped by whether ordinary statutory amendment can close them.

**Table 2: Seven Obligation Pairs**

| Group                    | EU anchor                                 | VN anchor                                     | ASEAN anchor                        | Key finding                                                                               |
|--------------------------|-------------------------------------------|-----------------------------------------------|-------------------------------------|-------------------------------------------------------------------------------------------|
| G1 Risk classification   | Art. 9, continuous, {Provider}            | Article 10(1)/(2)/(5), {Provider}→{Regulator} | Accountability principle, Voluntary | Article 10(5) closes the post-market void but surfaces a Provider-vs-Regulator mismatch   |
| G2 Conformity assessment | Art. 43(2), {Provider}                    | Article 13(2)(b), {Provider}                  | Not addressed                       | Total convergence on both strictness ( $\beta$ ) and obligor ( $\omega$ )                 |
| G3 Human oversight       | Art. 14, tier-gated, {Provider, Deployer} | Article 4(2), universal, $\omega = \emptyset$ | Human-centricity, Voluntary         | VN is stricter at Minimal risk but names no obligor — my headline reversal                |
| G4 Incident reporting    | Art. 12, {Provider, Deployer}             | Article 10(3) + Decree 142, {Provider}        | Disclosure, Recommended             | VN is stricter at Medium tier; reveals a proper-subset obligor mismatch                   |
| G5 Transparency          | Art. 50, {Provider}                       | Article 11(1), {Provider}                     | Transparency principle, Voluntary   | Total convergence — the only non-trivial structural alignment in the dataset              |
| G6 Risk-mgmt temporality | Art. 9, continuous, {Provider}            | Decree 142, event-triggered, {Deployer}       | Not addressed                       | Disjoint-actor mismatch; the direction of the strictness gap flips depending on risk tier |

| Group                      | EU anchor                | VN anchor                          | ASEAN anchor           | Key finding                                                             |
|----------------------------|--------------------------|------------------------------------|------------------------|-------------------------------------------------------------------------|
| G7 Risk taxonomy structure | Annex III, fixed catalog | Vietnam’s risk Catalog, multi-axis | 6 harm-type categories | Incomparable — the fundamental structures of the catalogs clash (§ 6.5) |

## 5. Quantitative Results

I ran all 2,880 possible contexts through my Python engine (evaluating the entire dataset in under a second).

**Table 3: Per-Group Gap Table**

| Group                                    | EU > VN            | VN > EU              | EU > ASEAN           | VN > ASEAN                                             |
|------------------------------------------|--------------------|----------------------|----------------------|--------------------------------------------------------|
| G1                                       | 240 (8.3%)         | 1,248 (43.3%)        | 1,152 (40.0%)        | 2,160 (75.0%)                                          |
| G2                                       | 576 (20.0%)        | 0 (0.0%)             | 1,152 (40.0%)        | 576 (20.0%)                                            |
| G3                                       | 0 (0.0%)           | 1,728 (60.0%)        | 1,152 (40.0%)        | 2,880 (100%)                                           |
| G4                                       | 432 (15.0%)        | 1,008 (35.0%)        | 1,152 (40.0%)        | 1,728 (60.0%)                                          |
| G5                                       | 0 (0.0%)           | 0 (0.0%)             | 1,440 (50.0%)        | 1,440 (50.0%)                                          |
| G6                                       | 576 (20.0%)        | 864 (30.0%)          | 1,152 (40.0%)        | 1,440 (50.0%)                                          |
| <b>Union (<math>\geq 1</math> group)</b> | <b>864 (30.0%)</b> | <b>1,728 (60.0%)</b> | <b>2,016 (70.0%)</b> | <b>2,880 (100%) / 2,700 (93.8%) excl. Article 4(2)</b> |

**Table 4: ObligorGap Incidence**

| Group        | $\omega_{EU}$        | $\omega_{VN}$ | ObligorGap   | % of total   | Type                                                       |
|--------------|----------------------|---------------|--------------|--------------|------------------------------------------------------------|
| G1           | {Provider}           | {Regulator}   | 576          | 20.0%        | Disjoint                                                   |
| G2           | {Provider}           | {Provider}    | 0            | 0.0%         | single-actor                                               |
| G3           | {Provider, Deployer} | $\emptyset$   | <b>1,152</b> | <b>40.0%</b> | Convergence                                                |
| G4           | {Provider, Deployer} | {Provider}    | 144          | 5.0%         | Named-set vs. unnamed                                      |
| G5           | {Provider}           | {Provider}    | 0            | 0.0%         | Proper-subset                                              |
| G6           | {Provider}           | {Deployer}    | 576          | 20.0%        | Convergence                                                |
| <b>Union</b> | —                    | —             | <b>1,152</b> | <b>40.0%</b> | Disjoint single-actor = G3 alone (G1/G4/G6 $\subseteq$ G3) |

*A note on the union calculations:* Simply adding up the per-group ObligorGap percentages gives a misleading 85%. In reality, groups 1, 4, and 6 are strictly nested inside group 3 (all are confined to High/Unacceptable risk tiers). The correct distinct-context union is exactly 1,152 contexts, or 40% of the dataset. Interestingly, ObligorGap (40%) and bindingness gaps (30–60%) are anti-correlated, meaning neither issue dominates the other completely.

*A note on EU consistency:* Five out of six EU rules I tracked trigger identically when `risk_tier`  $\in$  {High, Unacceptable}. Almost all the structural diversity in this dataset comes from the Vietnam and ASEAN frameworks.

## 6. Key Empirical Findings

**6.1 The Reverse Gap at Minimal Risk.** This is my principal finding: Vietnam’s Article 4(2) enforces binding rules at the Minimal risk tier where the EU’s Article 14 remains completely silent. I found a similar inversion in Group 4, where Vietnam’s Article 10(3) triggers at the Medium tier, capturing a broader scope than the EU’s High/Unacceptable gate. *Caveat:* “Stricter” here refers purely to legal force ( $\beta$ ). While Vietnam mandates a broad outcome, the EU often specifies concrete design mandates. Whether broad outcomes or specific mechanisms carry more practical weight is a separate debate.

**6.2 Vietnam vs. ASEAN Divergence and Transparency Alignment.** Vietnam diverges from the ASEAN baseline in nearly 100% of contexts (93.8% if I exclude Article 4(2)). Given how recent both documents are, this is uncharted territory for the region. On the flip side, Group 5 (Transparency) shows the exact opposite: complete alignment. The EU and Vietnam are both binding with identical support scopes. My Theorem C confirms this is the *only* non-trivial structural alignment across all twelve regime pairs.

**6.3 Arbitrage Direction Inversion.** Under my baseline encoding, an interesting dynamic emerges: the EU is the laxer jurisdiction twice as often as Vietnam (30% vs 60%). If a company were purely optimizing for regulatory arbitrage based on bindingness alone, it would more frequently find the EU to be the lighter-touch regime.

**6.4 Sensitivity Analysis.** I ran a  $2 \times 2$  sensitivity analysis to ensure these findings weren’t just artifacts of how I mapped the data. Four core findings held up under every test: the post-market void closure, the minimal-risk reversal, the VN $\leftrightarrow$ ASEAN divergence, and the transparency alignment. However, the arbitrage ratio (Finding 6.3) did *not* hold — under an alternative encoding, the EU became stricter than Vietnam (50% vs 40%). The takeaway here is that determining which regime is “laxer” is highly dependent on specific configurations, a nuance that prose-based analysis simply cannot capture.

**6.5 Failure Mode 1 — Classification-Scheme Incompatibility.** The EU categorizes risk using a static catalog of use-cases (Annex III), whereas Vietnam uses a multi-axis system based on harm scale, reversibility, and human approval. My analysis (Proposition H7.1 and Theorem C) proves that these two systems partition risk along irreconcilable axes. Neither framework naturally refines the other. (*Note: This assumes Vietnam’s draft Official Dispatch No. 1101/BKHCN-CNS&CDS remains structurally similar when finalized.*)

**6.6 Failure Mode 2 — Obligor Mismatches.** This is arguably the most operationally dangerous failure mode: regimes agree on a rule’s strictness but name different responsible parties. In a companion technical analysis, I proved the **Antichain Indeterminacy Theorem**, which shows that whenever co-firing rules form an antichain — no rule’s firing condition is a subset of another’s — standard legal aggregation (*lex specialis*) breaks down entirely and returns no answer, making union aggregation the only mathematically valid choice in that case. Vietnam’s Group 1 rules are exactly this case: every pairwise relation among them is incomparable, so all 816 of Group 1’s multi-rule contexts (out of 2,880) form antichains. Group 4 is the contrasting case — under a weaker syntactic specificity test, it is *not* an antichain, so *lex specialis* does return an answer there, but one that disagrees with union aggregation on 288 contexts. The two groups together show that aggregation choice can manufacture or dissolve an obligor gap depending on the rules’ structure, which is checkable by inspection before running any computation.

I confirmed four distinct ObligorGap instances affecting 40% of the scenarios. For example, the

EU assigns human oversight to {Provider, Deployer}, while Vietnam states a universal principle but names no specific actor ( $\emptyset$ ). In other cases, they target completely disjoint actors (e.g., Provider vs. Regulator, or Provider vs. Deployer).

## 7. Discussion and Limitations

**Legal force vs. operational specificity.** To be clear, I am measuring formal legal force ( $\beta$ ), not on-the-ground enforcement capability. Vietnam’s approach mandates a binding outcome; the EU mandates an outcome plus a specific mechanism. I am only quantifying the former.

**Limitations.** 1. *Shared schema:* I merged the EU’s four-level and Vietnam’s three-level risk tiers based on reasonable assumptions. While the directional arbitrage ratio is sensitive to this choice, my four structural findings are not. 2. *Interpretation:* Coding laws into logic inherently requires human interpretation. My theorems prove properties of the model, not the absolute legal reality of the statutes. 3. *Single-researcher verification:* I conducted the statutory research and obligor annotations independently, without an audit by a second legal expert. The code is open and reproducible, but the inputs should be viewed with this in mind.

**Dual-Use Considerations.** A tool that pinpoints the exact scenarios where two laws differ is essentially an arbitrage detector. It can help policymakers close loopholes, but it can also help companies exploit them. I chose to publish this because well-resourced legal teams are already doing this analysis manually; releasing an open, computable tool levels the playing field rather than hiding the asymmetry.

**Future Work.** I see two immediate next steps: expanding the context space to model multi-actor, cross-border incident chains, and fully characterizing the conditions under which *lex specialis* can successfully resolve regulatory overlaps.

## 8. Conclusion and Recommendations

Vietnam’s AI Law is architecturally distinct from the EU AI Act, frequently stricter in formal terms, and significantly ahead of ASEAN’s voluntary baseline. Two of the four major harmonization failures I found — partition incompatibility and obligor mismatch — cannot be resolved by simply editing text; they require foundational architectural decisions.

Because I built this as an executable model, anyone can plug their specific AI system’s parameters into the HSDL engine to check compliance across these three jurisdictions in seconds, bypassing weeks of comparative legal research.

1. **For ASEAN:** The bloc needs to decide on its organizing risk architecture (tier-gated vs. multi-axis) *before* drafting binding DEFA rules. Vietnam’s law proves that emerging economies do not have to default to the EU’s design. Furthermore, if mutual recognition is the goal, laws must explicitly allocate responsibility (obligors) — agreeing on the strictness of a rule is useless if countries disagree on who must comply.
2. **For Vietnam:** Policymakers should explicitly state whether the strictness at the minimal-risk tier is a deliberate choice or an oversight. If they want to converge with the EU, they actually need to *loosen* Article 4(2). Vietnam also needs to finalize the Article 13(4) Catalog quickly, as it dictates how third-party certifications will function.

3. **For Deployers:** You cannot ensure cross-border compliance simply by following the “strictest” regime. Because 40% of the scenarios involve an obligor mismatch, assuming the EU will cover you in Vietnam (or vice versa) is a systemic risk. I recommend running your specific configuration through the HSDL tool rather than relying on high-level country intuitions.

## Code and Data Availability

- **Code:** All verification scripts (`phase3_omega.py`, `verify_g4_lex.py`, `theorem_c_matrix.py`, `verify_union.py`, and the HSDL reference implementation) are available under the Apache 2.0 license at: <https://github.com/Eilodon/HSDL-Harmonization-Gap>.
- **Data:** There are no proprietary datasets. All evaluated statutes (the EU AI Act, Vietnam’s AI Law No. 134/2025/QH15, the ASEAN AI Governance Guide, and Decree No. 142/2026/NĐ-CP) are public.
- **Supplementary outputs:** Two companion analyses (a deep dive into *lex specialis* math and a plain-language policy translation) are available upon request.

## Author Contributions

As the sole author, I led the project, developed the HSDL formal framework and its extension into the legal-harmonization domain, conducted the statutory research and obligor annotation, implemented and ran the verification scripts, and wrote the manuscript.

## LLM Usage Statement

I used Large Language Models during this project to synthesize background literature, draft experimental verification code, condense a longer independent manuscript into this format, and perform stylistic editing. However, the formal framework (Definitions H1–H13, Theorems A–C), the statutory research, the rule encoding, and all quantitative results were produced and verified entirely by me, independently of LLM generation. I strictly reviewed and tested all LLM-assisted code drafts before using them as evidence. All results are fully reproducible by re-running my scripts against the cited legal texts.

## References

1. Boella, G., Governatori, G., Rotolo, A., & van der Torre, L. (2010). “Lex Minus Dixit Quam Voluit, Lex Magis Dixit Quam Voluit: A Formal Study on Legal Compliance and Interpretation.” In Casanovas, P., Pagallo, U., Sartor, G., & Ajani, G. (eds.), *AI Approaches to the Complexity of Legal Systems* (AICOL 2009), LNCS vol. 6237, pp. 162–183. Springer.
2. Governatori, G., & Rotolo, A. (2010). “Changing legal systems: legal abrogations and annulments in Defeasible Logic.” *Logic Journal of the IGPL*, 18(1), 157–194.
3. Guldemann, T., et al. (2024). “COMPL-AI: A Computational Framework and Benchmark Suite for EU AI Act Compliance.” *arXiv preprint*.
4. Lu, Y., & Tie, F. H. (2025). “A comparative analysis of artificial intelligence regulation in

- ASEAN and the European Union.” *Journal of Governance and Regulation*, 14(4).
5. Marwala, T. (2026). “The AI Governance Arbitrage.” UN University Working Paper.
  6. Palmirani, M., et al. (2011). “LegalRuleML: XML-Based Rules and Norms.” *RuleML 2011*, LNCS 7018, Springer.
  7. Modern Diplomacy (2026). “From Soft Law to Hard Rules: Pushing for Binding AI Governance in ASEAN.” January 2026.
  8. ASEAN Secretariat. “ASEAN Digital Economy Framework Agreement” — negotiations concluded May 2026; signature expected November 2026.
  9. Vietnam’s AI Law (Law No. 134/2025/QH15) & Decree No. 142/2026/NĐ-CP.