From Living Within a Lie to Architectural Truth

Legitimacy Rupture and the Re-Design of Middle-Power Cooperation

Executive Thesis

Mark Carney's Davos speech marks more than a shift in tone. It represents a rare public acknowledgment by a system-level insider that global institutional legitimacy has crossed a **non-recoverable threshold**. The rules-based order has not merely weakened; it has become **performative**—maintained through ritual language that no longer corresponds to lived institutional reality.

This paper argues that the present moment is not a transition between orders, but a **legitimacy rupture**. Attempts to preserve cooperation by invoking obsolete legitimacy narratives now actively suppress institutional learning, distort strategic choice, and accelerate breakdown dynamics—particularly for middle powers.

Using the Institute for Regenerative Systems Architecture (IRSA) framework, this paper diagnoses the rupture and outlines how cooperation can be **re-architected without pretending**—through truth-aligned legitimacy, explicit operating architectures, and credible commitment mechanisms that do not rely on hegemonic enforcement or shared fictions.

1. The Rupture, Not the Transition

Periods of institutional stress are often described as transitions. The language is comforting. It implies continuity, recoverability, and an eventual return to stability. But not all breakdowns are transitional. Some are **terminal to the legitimacy structure that preceded them**.

The present moment belongs to the latter category.

This paper argues that the global institutional order is not undergoing a difficult passage between stable equilibria, but has crossed a **legitimacy rupture**—a threshold beyond which the narratives that once sustained cooperation can no longer be credibly restored.

1.1 Why "Rules-Based Order" Language Now Suppresses Learning

For decades, the phrase "rules-based order" served an important function. It coordinated expectations, disciplined behaviour, and legitimised constraint by embedding power within shared norms.

Today, that language no longer performs those functions reliably.

Instead, it increasingly acts as a **learning suppressant**.

When institutions invoke rules that are selectively enforced, asymmetrically applied, or routinely overridden, the language ceases to clarify reality. It obscures it. Actors learn that outcomes depend not on rules as stated, but on exceptions, leverage, and context.

Yet because the language persists, institutions are discouraged from redesigning themselves around what actually governs behaviour. Learning stalls at precisely the moment it is most needed.

1.2 Breakdown vs Transition: A Structural Distinction

A transition implies that:

- legitimacy remains broadly intact,
- institutional learning channels remain open,
- and reform operates within the same narrative frame.

A rupture is different.

A rupture occurs when:

- institutional claims no longer correspond to operational reality,
- learning signals are systematically filtered or ignored,
- and reform proposals must contradict the institution's own self-description to be effective.

At this point, continuity becomes performative. Reform efforts fail not because they are insufficiently ambitious, but because they are framed inside a legitimacy story that no longer holds.

1.3 Narrative Lag as a Source of Systemic Risk

Narrative lag refers to the persistence of legitimacy language after its enabling conditions have disappeared.

This lag is not benign. It introduces three compounding risks:

1. **Misdiagnosis**
   Problems are framed as temporary deviations rather than structural failures.

2. **Delayed Adaptation**
   Institutions invest in repair strategies instead of redesign.
3. **Trust Drain**
   Participants recognise the mismatch between words and outcomes long before institutions do.

By the time legitimacy is publicly questioned, it has usually already collapsed privately.

1.4 The LGIT Threshold: When Reaffirmation Stops Working

Within the LGIT framework, legitimacy erosion follows a predictable pattern:

- early strain produces calls for recommitment,
- mid-stage decay produces procedural reform,
- late-stage decay produces symbolic escalation.

Rupture occurs when reaffirmation itself becomes destabilising.

At this point:

- repeating the legitimacy narrative increases grievance,
- moral appeals accelerate withdrawal,
- and institutional insistence is interpreted as denial.

This is the threshold the current system has crossed.

1.5 Why Acknowledgment Matters More Than Authority

What makes recent acknowledgments significant is not who speaks, but **what is admitted**.

When system insiders publicly recognise that the old order no longer constrains behaviour, they are not proposing an alternative. They are removing a sign that no longer corresponds to the terrain.

This matters because institutions cannot be redesigned while pretending that they still operate as described.

Acknowledgment does not solve the problem—but it makes **architecture possible**.

1.6 The Cost of Misnaming the Moment

Misnaming rupture as transition has real consequences.

It leads institutions to:

- double down on compliance rituals,
- escalate moral rhetoric,

- and exhaust political capital defending indefensible claims.

Most dangerously, it leaves middle powers without a language to explain why existing arrangements no longer work—forcing them to choose between silent performance and reputational rupture.

Naming rupture is not pessimism. It is **precondition for design**.

1.7 Clearing the Ground for Architecture

Rupture does not mean cooperation ends. It means cooperation must be **re-founded**.

Once obsolete legitimacy narratives are removed, institutions can:

- scope commitments honestly,
- reopen learning channels,
- design enforcement pathways,
- and rebuild legitimacy from alignment rather than assertion.

This is not regression. It is maturation under constraint.

The question is no longer whether the old order can be restored. It is whether new architectures can be built **before grievance hardens into disengagement**.

2. Living Within a Lie as Institutional Strategy

When legitimacy decays beyond repair, institutions rarely collapse immediately. Instead, they adapt in a more subtle—and ultimately more corrosive—way: they continue to operate **as if** their legitimacy remains intact.

This condition can be described as *living within a lie*—not as a moral indictment, but as an institutional survival strategy under constraint.

In this phase, institutions do not deny reality outright. They bracket it. They preserve outward coherence by separating what is said from what is known, and what is known from what can be acted upon.

2.1 Ritual Compliance as a Substitute for Legitimacy

In healthy institutions, compliance is a byproduct of legitimacy. In late-stage institutions, compliance becomes a **proxy** for legitimacy.

Ritual compliance refers to the continued performance of institutional behaviours—meetings, communiqués, reporting cycles, affirmations of shared values—after those behaviours have ceased to meaningfully constrain action or guide outcomes.

These rituals persist because they perform three stabilising functions:

- They signal continued membership in a recognised order
- They reduce coordination risk by preserving shared expectations
- They delay the reputational cost of open divergence

However, ritual compliance is not neutral. It actively suppresses adaptation by rewarding those who perform alignment rather than those who surface misalignment.

Over time, institutions begin to optimise for *appearing legitimate* rather than *being operationally truthful*.

2.2 Performance vs Truth-Aligned Participation

As ritual compliance becomes dominant, participation bifurcates.

On one side is **performance-based participation**:

- Actors affirm principles they privately discount
- Deviations are hidden rather than addressed
- Decision rationales are retrofitted to match accepted narratives

On the other side is **truth-aligned participation**:

- Actors describe constraints as they are
- Commitments are scoped to enforcement capacity
- Trade-offs are acknowledged rather than moralised

Late-stage institutions implicitly punish truth-aligned participation. Those who surface limits, asymmetries, or enforcement gaps are seen as destabilising—even when they are descriptively accurate.

The result is a self-reinforcing dynamic: the more institutions rely on performance, the more reality must be suppressed; the more reality is suppressed, the more performance is required.

2.3 The LGIT Loop: How Lies Compound Grievance

Within the LGIT framework, living within a lie accelerates grievance accumulation in three ways.

First, it **invalidates experience**. When institutional narratives no longer match lived outcomes, participants internalise the belief that their perceptions are either irrelevant or dangerous to express.

Second, it **blocks grievance articulation**. Because the institution cannot acknowledge structural limits, grievances are reframed as behavioural failures, misunderstandings, or bad faith.

Third, it **misattributes blame**. Systemic breakdown is personalised, producing scapegoating, moral fatigue, and disengagement.

Crucially, grievance does not disappear when it cannot be named. It migrates—into withdrawal, parallel coordination, or sudden rupture when accumulated pressure exceeds tolerance.

2.4 Why Middle Powers Are Structurally Trapped

Middle powers are disproportionately exposed to this dynamic because they sit at the intersection of expectation and constraint.

They are:

- Expected to uphold institutional norms
- Incentivised to demonstrate reliability
- Limited in their capacity to enforce reciprocity

Living within the lie becomes particularly costly for middle powers. Continued performance sustains a system they cannot steer, while open divergence risks exclusion from coordination mechanisms they still depend on.

This creates a structural trap:

- Perform, and absorb legitimacy risk without agency
- Defect, and bear reputational cost without alternatives

IRSA theory suggests that this is not a failure of courage or clarity. It is an architectural deadlock.

2.5 From Survival Strategy to Failure Mode

What begins as a survival strategy eventually becomes a failure mode.

Institutions living within a lie:

- Lose the ability to distinguish signal from noise
- Become reactive rather than adaptive
- Interpret realism as betrayal

At this stage, collapse often appears sudden—but it is not. It is the delayed consequence of long-suppressed learning.

The central danger is not that institutions acknowledge reality too soon, but that they do so **too late**, after trust has already drained away.

2.6 The Architectural Exit: Designing for Truth Without Collapse

The alternative to living within a lie is not reckless disclosure or moral purity. It is **architectural truthfulness**.

Architectural truthfulness means:

- designing institutions that can acknowledge constraint without losing coherence,
- enabling scoped participation rather than universal pretence,
- and embedding learning and revision as normal operations rather than crisis responses.

For middle powers, this is the critical pivot. Truth-aligned institutions are not weaker—they are **more resilient**, because they fail smaller, adapt faster, and retain legitimacy precisely by not over-claiming it.

This sets the stage for the next question: if universal legitimacy is no longer viable, how should cooperation be structured?

That question leads directly to the problem of **variable geometry**.

3. Strategic Autonomy vs Fortress Logic

One of the most important—and easily overlooked—moves in recent system-level discourse is the rejection of the supposed choice between *strategic autonomy* and *fortress logic*. This dichotomy has quietly shaped much middle-power thinking over the past decade, yet it is architecturally incoherent.

The dichotomy assumes that sovereignty must either be pooled within a universal rules-based system or reclaimed through insulation and exclusion. IRSA theory treats this as a category error: it confuses *control over commitment* with *control over territory*.

3.1 Why the Dichotomy Persists Despite Its Failure

The persistence of this framing is not accidental. Each pole offers a psychologically and politically convenient response to legitimacy decay.

Strategic autonomy rhetoric persists because it:

- preserves symbolic sovereignty without confronting enforcement design,
- allows institutions to speak in universal terms while acting selectively,
- and postpones the costs of architectural redesign.

Fortress logic persists because it:

- offers visible control under grievance pressure,
- substitutes exclusion for institutional coherence,
- and creates the appearance of decisiveness when legitimacy is fragile.

Both responses are understandable. Both are insufficient.

In practice, strategic autonomy without enforcement capacity produces *aspirational sovereignty*, while fortress logic produces *brittle isolation*. Neither addresses the core problem: how commitments are made credible, revised, or exited under constraint.

3.2 Sovereignty Without Enforcement Is Symbolic

Sovereignty is not defined by the ability to declare intent, but by the ability to uphold commitments.

Where enforcement pathways are unclear, autonomy becomes symbolic rather than operational. Actors may retain formal discretion, but lose practical control as coordination failures, reputational exposure, and escalation risks accumulate.

Conversely, fortress logic mistakes insulation for resilience. By narrowing exposure rather than designing it, institutions reduce learning capacity and increase the cost of inevitable re-engagement.

Both approaches attempt to *avoid architecture*. Both fail for the same reason.

3.3 The Architectural Resolution: Autonomy Through Design

IRSA dissolves the autonomy–fortress dichotomy by relocating sovereignty at the level of **commitment architecture**, not narrative posture.

Under this framing:

- autonomy means control over *when* commitments bind and *how* they evolve,
- resilience means *designed exposure* rather than withdrawal,
- and cooperation is conditional, legible, and enforceable by design.

Commitment & Enforcement Architecture (CEA) governs credibility: how deviations are detected, addressed, and escalated without rupture.
 Institutional Operating Architecture (IOA) governs coherence: how authority is exercised, revised, and preserved across time.

Together, these architectures enable institutions to remain open without being naïve, and autonomous without becoming isolated.

This is not a third option between autonomy and fortress logic. It is a reframing that renders the dichotomy obsolete.

4. Variable Geometry as Proto-Architecture

As universal legitimacy erodes, institutions instinctively reach for flexibility. Cooperation becomes conditional, selective, and context-dependent. This is the logic behind what is often described as *variable geometry*.

Variable geometry acknowledges a basic truth: not all actors can, or should, commit to the same obligations at the same time. In this sense, it represents an important break from universalist pretence. However, flexibility alone does not constitute architecture.

Without design discipline, variable geometry remains a **proto-architecture**—directionally correct, but structurally incomplete.

4.1 Why Uniform Multilateralism Is Exhausted

Uniform multilateralism assumes that:

- commitments are symmetrical,
- enforcement is implicit,
- and legitimacy is shared by default.

These assumptions no longer hold. Enforcement capacity is uneven, incentives diverge, and legitimacy is fractured across constituencies. Maintaining uniformity under these conditions does not preserve cooperation; it accelerates disengagement.

Variable geometry emerges as a pragmatic response:

- smaller coalitions,
- differentiated commitments,
- and modular participation.

This is not fragmentation—it is **realism**. But realism without structure is unstable.

4.2 What Variable Geometry Gets Right

Variable geometry correctly recognises that:

- cooperation must be scoped rather than universal,
- legitimacy must be earned locally rather than assumed globally,
- and institutions must tolerate difference without collapse.

It creates space for:

- middle-power alignment without hegemonic sponsorship,
- experimentation without total system buy-in,
- and cooperation among actors with asymmetric capacities.

In short, it restores *possibility* where uniformity had become paralyzing.

4.3 Where Variable Geometry Fails Without Architecture

Despite its appeal, variable geometry often fails in practice because it lacks **operability**.

Common failure modes include:

- ad hoc coalitions with no shared learning pathways,
- commitments that are politically announced but operationally vague,
- and cooperation that depends on personal relationships rather than institutional design.

In these cases, flexibility becomes fragility. Without explicit architecture, variable geometry:

- multiplies interfaces without coordination,
- diffuses accountability,
- and increases transaction costs over time.

What begins as adaptive cooperation gradually degrades into managed incoherence.

4.4 Flexibility vs Operability

The core distinction is between **flexibility** and **operability**.

- Flexibility allows variation
- Operability sustains coordination under variation

Institutions can be flexible and still fail if participants do not know:

- how decisions are made,
- how commitments are revised,
- how deviations are handled,
- or how learning is shared across modules.

Operability requires architecture. It cannot be improvised.

4.5 Variable Geometry Needs an Operating Layer

For variable geometry to mature from proto-architecture into durable system design, it must be paired with an explicit **operating layer**.

This operating layer specifies:

- participation rules across different cooperation modules,
- interfaces between aligned but non-identical institutions,
- and escalation pathways when coordination breaks down.

Without this layer, variable geometry depends on continual renegotiation—an approach that exhausts political capital and erodes trust.

With it, differentiated cooperation can **compound** rather than fragment.

4.6 From Strategy to System

The critical mistake is treating variable geometry as a *strategy* rather than a *system*.

Strategies respond to conditions. Systems endure through them.

To function over time, variable geometry must be:

- legible to participants,
- predictable under stress,
- and capable of learning from partial failure.

These properties do not emerge organically. They must be designed.

This is where Institutional Operating Architecture becomes decisive. IOA transforms variable geometry from a pragmatic workaround into a coherent, scalable form of post-hegemonic cooperation.

4.7 Preparing the Ground for Re-Architected Legitimacy

Variable geometry clears the ground by abandoning universal pretence. Operating architecture determines what can be built next.

Without IOA, flexibility postpones rupture.
 With IOA, flexibility becomes resilience.

This prepares the ground for the central claim of this paper: legitimacy in the next order will not be universal or assumed. It will be **architected, scoped, and earned**.

5. Re-Architecting Legitimacy (IRSA Framework)

Legitimacy cannot be restored by declaration, reaffirmation, or rhetorical escalation. Once institutions cross a rupture threshold, legitimacy becomes a **design property**, not a communicative one.

IRSA treats legitimacy as an *emergent outcome* of alignment across institutional layers. When those layers drift out of sync, legitimacy decays regardless of intent. When they are re-aligned, legitimacy can re-emerge—even under constraint.

This section outlines how legitimacy can be re-architected through four mutually reinforcing layers: LGIT, ILA, CEA, and IOA.

5.1 LGIT: From Performed Legitimacy to Truth-Aligned Legitimacy

In late-stage institutional decay, legitimacy is often mistaken for *compliance*. IRSA's Legitimacy–Grievance–Institutional Trajectory (LGIT) framework rejects this equivalence.

Legitimacy arises when participants believe that:

- the institution accurately describes reality,
- grievances are recognisable and actionable,
- and participation has a plausible causal relationship to outcomes.

When institutions continue to assert universality while operating selectively, they generate **legitimacy dissonance**. Participants comply outwardly while withdrawing inwardly. Grievance does not disappear—it compounds silently.

Truth-aligned legitimacy does not require moral consensus or universal agreement. It requires **descriptive honesty**:

- naming enforcement asymmetries,
- acknowledging bounded commitments,
- and scoping obligations to what can actually be upheld.

Paradoxically, institutions that admit limitation often regain trust faster than those that insist on totality. Legitimacy begins to recover not when institutions claim to be fair, but when they stop pretending to be something they are not.

5.2 ILA: Reopening Institutional Learning After Narrative Lock-In

Once legitimacy becomes performative, institutional learning shuts down.

Institutional Learning Architecture (ILA) explains why. Learning depends on feedback being:

- observable,
- admissible,
- and consequential.

In narrative-locked institutions, feedback becomes reputationally dangerous. Signals that contradict the dominant legitimacy story are reframed as threats rather than inputs. This produces **anti-learning regimes**: environments where adaptation is structurally discouraged.

Reopening learning requires architectural intervention, not cultural exhortation.

ILA re-designs institutions so that:

- deviation is expected rather than denied,
- error is informative rather than embarrassing,
- and revision is procedural rather than exceptional.

For middle-power cooperation, this is critical. Institutions must be able to learn *in public* without collapsing legitimacy. This is only possible when legitimacy is grounded in truth-alignment rather than perfection.

5.3 CEA: Credible Commitment Without Hegemonic Enforcement

The failure of global cooperation is often attributed to declining norms. IRSA identifies a more precise cause: **commitment without credible enforcement**.

Commitment & Enforcement Architecture (CEA) distinguishes between:

- declared commitments,
- perceived commitments,
- and enforceable commitments.

In hegemonic systems, enforcement is externalised. In post-hegemonic conditions, enforcement must be **designed**.

CEA enables cooperation by making explicit:

- how commitments are monitored,
- how deviations are addressed,
- how proportional consequences are applied,
- and how escalation occurs without rupture.

For middle powers, this is not optional. Without CEA, cooperation relies on moral pressure and reputational risk—both of which collapse under strain. With CEA, cooperation becomes conditional but durable.

Importantly, CEA does not require coercion. It requires **predictability**. Actors will accept constraint when enforcement pathways are visible, bounded, and symmetrical.

5.4 IOA: Ensuring Institutions Operate as Described

Even with aligned legitimacy, learning, and commitment, institutions fail if they do not operate as their own descriptions suggest.

Institutional Operating Architecture (IOA) governs the *actual mechanics* of participation:

- who decides what, when, and how,
- how authority accumulates or dissipates,
- how procedures adapt over time,
- and how institutional memory is preserved.

Many contemporary institutions suffer from **operational drift**. Their formal mandates remain intact, but their day-to-day functioning evolves informally in response to pressure. Over time, the gap between description and operation becomes a legitimacy fault line.

IOA closes this gap by:

- making operating assumptions explicit,
- aligning incentives with stated purpose,
- and embedding revision pathways into institutional design.

Without IOA, even well-intentioned reforms degrade into ritual compliance. With IOA, institutions can evolve without losing coherence.

5.5 Alignment Across Layers: Legitimacy as an Emergent Property

The central IRSA insight is that legitimacy cannot be fixed at a single level.

- LGIT without ILA produces brittle moralism
- ILA without CEA produces insight without consequence
- CEA without IOA produces enforcement without coherence
- IOA without truth-aligned legitimacy produces efficient emptiness

Legitimacy re-emerges only when all four layers are aligned.

This alignment does not restore the old order. It enables **post-fictional cooperation**: institutions that function without requiring participants to affirm what they no longer believe.

For middle powers, this is not a consolation prize. It is a strategic advantage.

6. The Opportunity for Middle Powers

Legitimacy does not scale from the top down. In post-hegemonic conditions, it compounds laterally.

This creates a distinct opportunity for middle powers—not because they are morally superior or strategically bolder, but because they are already operating under the constraints the system as a whole is moving toward.

6.1 Why Middle Powers Can Architect First

Middle powers are structurally positioned to lead institutional redesign because they already live in a post-fictional reality.

They typically operate with:

- bounded and selective commitments,

- limited enforcement capacity,
- high sensitivity to reputational risk,
- and strong incentives for stable cooperation without domination.

In other words, they already coordinate without universal legitimacy or hegemonic backing. What they lack is not realism, but **explicit architecture**.

IRSA does not ask middle powers to abandon existing institutions or values. It asks them to formalise what they are already doing implicitly—and to design for durability rather than improvisation.

6.2 How Legitimacy Compounds When Shared

When multiple middle powers adopt truth-aligned architectures, legitimacy becomes **networked rather than universal**.

Instead of a single order demanding compliance, legitimacy emerges from:

- interoperable institutions,
- shared commitment standards,
- and mutual recognition of constraints.

This form of legitimacy is more resilient precisely because it is partial. It does not collapse when one node fails, nor does it require consensus to function.

Crucially, legitimacy in this model is not asserted. It is *earned through alignment*.

6.3 What Middle-Power Institutions Look Like When Designed Honestly

Institutions built for post-rupture cooperation share common characteristics:

- They do not claim universality
- They publish their enforcement limits
- They allow modular, opt-in participation
- They include explicit revision and exit pathways
- They prioritise operability over symbolism

Such institutions do not promise stability. They promise *learnability*. They do not eliminate conflict. They make it governable.

6.4 From First Movers to Structural Anchors

Middle powers that move first do not replace the old order. They provide **anchors for the next one**.

By demonstrating that cooperation can function without pretence, they lower the cost for others to follow. Over time, these architectures become reference points—not because they are imposed, but because they work.

This is how legitimacy returns: not as a universal claim, but as a shared property of institutions that tell the truth and operate accordingly.

An Illustrative Middle-Power Cooperation Architecture (Not a Policy Proposal)

To make the architectural claims in this paper concrete, consider a *hypothetical* middle-power cooperation arrangement — for example, a climate-security or supply-chain resilience compact among a small group of middle powers.

This example is illustrative only. It is not a recommendation or blueprint.

Scope and Membership

- Participation is modular and opt-in, not universal.
- Members commit only to obligations they can credibly enforce domestically.
- Exit is explicit and non-punitive.

Commitment & Enforcement Architecture (CEA)

- Commitments are tiered rather than binary.
- Monitoring focuses on observable actions, not stated intent.
- Non-compliance triggers graduated responses (loss of privileges, not moral condemnation).
- Escalation pathways are defined in advance, reducing crisis bargaining.

Institutional Learning Architecture (ILA)

- Deviations are recorded as design signals, not treated as norm violations.
- Periodic review cycles are mandatory and expected.
- Revision is procedural, not exceptional.

Institutional Operating Architecture (IOA)

- Decision rights are explicit: who can revise commitments, pause participation, or trigger escalation.
- Authority is bounded and time-limited.
- Institutional memory is preserved independently of political cycles.

Legitimacy Outcome
Legitimacy in this arrangement does not derive from universal participation or moral claims. It emerges because:

- commitments are believable,
- enforcement is predictable,
- learning is visible, and
- institutions operate as described.

This is what post-rupture cooperation looks like: not idealised, but resilient.

7. Conclusion: From Naming Rupture to Designing Reality

This paper has argued that the present moment is not a difficult transition within a stable institutional order, but a legitimacy rupture that cannot be repaired through reaffirmation, escalation, or nostalgia. Once institutional claims cease to correspond to operational reality, legitimacy does not weaken gradually—it **changes category**. It becomes architectural.

Attempts to preserve cooperation by continuing to invoke obsolete legitimacy narratives do not stabilise the system. They suppress learning, compound grievance, and increase the probability of sudden breakdown. In this context, "living within a lie" is not merely unsustainable; it is actively destabilising.

The implication is not that cooperation must end, but that it must be **re-founded**.

IRSA's central contribution is to show that legitimacy does not reside in declarations, norms, or values alone. It emerges when institutions are designed so that:

- they describe reality as it is (LGIT),
- they can learn without collapsing trust (ILA),
- they make commitments credible without domination (CEA),
- and they operate as they claim to operate (IOA).

When these layers align, legitimacy reappears—not as a universal assertion, but as a property of systems that work.

This re-framing dissolves several false choices that have paralysed institutional response: between realism and cooperation, autonomy and openness, flexibility and stability. Architecture makes these trade-offs governable. Without it, they remain rhetorical and unresolved.

For middle powers, this moment is not primarily a crisis of relevance, but a **window of architectural agency**. Precisely because they cannot rely on hegemonic enforcement or universal legitimacy, they are best positioned to design institutions that function under

constraint. By formalising what has previously been improvised, middle powers can create cooperation that is scoped, legible, and resilient.

Taking down the sign that no longer matches the terrain is only the first step. What follows determines whether grievance hardens into disengagement or is channelled into institutional renewal.

The next order will not be restored. It will be built—by institutions willing to tell the truth about power, commitment, and limitation, and to design cooperation that holds because it aligns with reality, not because it invokes it.

Diagram 1: *Legitimacy Rupture vs Transition*

Where: Section 1
 Purpose: Make rupture visually undeniable

Axes:

- X: Time

- Y: Legitimacy alignment (claims vs reality)

Shows:

- Transition = dip + recovery

- Rupture = divergence + narrative lag

This becomes one of your canonical LGIT visuals.

Diagram 2: *Living Within a Lie (LGIT Loop)*

Where: Section 2
 Purpose: Show grievance compounding structurally

Cycle:
 Narrative → Ritual Compliance → Suppressed Learning → Misattributed Blame → Grievance → Withdrawal → (back to Narrative)

This is an IRSA signature diagram.

Diagram 3: *Strategic Autonomy vs Fortress Logic (False Dichotomy)*

Where: Section 3

Triangle or 2×2 showing:

- Symbolic sovereignty

- Fortress isolation

- Architectural autonomy (IRSA solution)

Visually shows that both poles fail for the same reason: no enforcement design.

Diagram 4: *Variable Geometry: Strategy vs System*

Where: Section 4

Two panels:

- Left: ad hoc variable geometry (fragile)

- Right: IOA-enabled variable geometry (compounding)

This diagram will be extremely reusable.

Diagram 5: The IRSA Legitimacy Stack

Legitimacy as an Emergent Property of Aligned Institutional Layers

---

1. Purpose of the Diagram

This diagram visually demonstrates the **core IRSA claim**:

> Legitimacy is not a narrative, norm, or value.
> It is an *emergent property* of alignment across institutional architectures.

It shows why legitimacy collapses when institutions "do the right things" at the wrong layer—or skip layers entirely.

This diagram turns Section 5 from theory into **mechanism**.

---

2. Diagram Structure (What to Draw)

Overall Form

- A **vertical stack of four layers**

- Each layer is a **distinct architectural system**

- Legitimacy is **not a layer** — it *emerges* when all layers align

Think: **systems stack, not hierarchy of importance**.

Layer 1 (Bottom): LGIT

Legitimacy–Grievance–Institutional Trajectory

Label:

LGIT — Truth-Aligned Legitimacy

Function:

- Aligns institutional claims with lived reality

- Determines whether participation feels meaningful or performative

Key properties (small text in layer):

- Descriptive honesty

- Grievance recognisability

- Plausible causal link between action and outcome

Layer 2: ILA

Institutional Learning Architecture

Label:

ILA — Learning Under Constraint

Function:

- Determines whether institutions can adapt without collapsing legitimacy

- Governs how feedback enters design

Key properties:

- Feedback admissibility

- Error visibility

- Procedural revision

Layer 3: CEA

Commitment & Enforcement Architecture

Label:

CEA — Credible Commitment

Function:

- Governs how commitments bind, weaken, or escalate

- Makes cooperation predictable rather than moralised

Key properties:

- Monitoring pathways

- Proportional consequence

- Escalation without rupture

Layer 4 (Top): IOA

Institutional Operating Architecture

Label:

> IOA — Operating as Described

Function:

- Governs how authority, decision rights, and procedures actually operate

- Prevents drift between mandate and behaviour

Key properties:

- Authority pathways

- Decision ownership

- Institutional memory

Emergent Property (Not a Layer)

Above the stack, floating or bracketed:

> **LEGITIMACY (Emergent)**
> *Appears only when all layers are aligned*

This is crucial: legitimacy is **shown as a result**, not a foundation.

---

3. Failure-Mode Annotations (This Is the Killer Feature)

On the **right-hand side**, annotate what happens when layers are missing or misaligned.

Use arrows or callouts.

Missing / Weak LGIT

> *"Performative legitimacy"*
> Institutions sound coherent but feel false.

LGIT + No ILA

"Brittle moralism"
Truth is asserted, but learning is punished.

ILA + No CEA

"Insight without consequence"
Institutions understand problems they cannot act on.

CEA + No IOA

"Enforcement without coherence"
Rules exist, but authority is fragmented.

IOA + No Truth Alignment

"Efficient emptiness"
Institutions operate smoothly toward meaningless or distrusted ends.

This is where many modern institutions sit.

---

4. Visual Style Guidance (Important)

- **Monochrome or muted palette** (no bright colours)

- Clean, architectural typography

- No arrows looping endlessly — this is *stack alignment*, not a cycle

- Avoid "pyramids" (this is not a maturity model)

This should feel closer to:

- systems engineering

- institutional design
  than management consulting.

---

5. How to Reference It in the Paper

In **Section 5**, add one sentence like:

> *Figure X illustrates legitimacy not as a declared attribute, but as an emergent property arising from alignment across four institutional architectures: LGIT, ILA, CEA, and IOA.*

That's all. Let the diagram do the work.

---

6. Why This Diagram Matters Strategically

This diagram:

- Differentiates IRSA from governance, ethics, and risk frameworks

- Explains *why* "doing more transparency" or "reaffirming norms" fails

- Makes legitimacy *designable* rather than moralised

- Is reusable across **every** IRSA domain (AI, libraries, philanthropy, geopolitics)

It will quietly become one of your **signature visuals**.