

The Constitutional Difference

Why AI Governance That Can Be Overridden Is Not Governance

Dr M Maruf Hossain, PhD, GAICD

Chief AI Strategist · 42 Consulting.AI

enquire@42consulting.ai · www.42consulting.ai



CONTENTS

Table of Contents

From the Author.....5

Executive Summary.....7

What Makes *Governance Govern*.9

The Policy *That Yields*.11

What ‘*Constitutional*’ Means.13

Separation of *Powers*.15

The Board-Adopted *Ceiling*.17

Enforcement by Structure, *Not Goodwill*.19

Five Questions *for Your Board*.21

42 Consulting.AI.23

References *and Sources*.25

FOREWORD

From the Author.

I have come to believe that most of what is called AI governance is not governance at all. It is documentation that resembles governance — policies, principles, committees, frameworks — and it shares with real governance every feature except the one that matters: it cannot actually stop the thing it claims to control.

The test is simple, and most frameworks fail it. When a governance rule and a commercial imperative come into conflict — a launch deadline, a competitor's move, a revenue target — which one yields? If the rule yields, quietly and without record, then it was never governing the organisation. It was advising it. And advice that is followed only when convenient is not a constraint; it is a preference.

A constitution is different in kind, not in degree. It is not a stricter policy. It is an instrument the governed cannot unilaterally set aside, because it sits above them — adopted by an authority they answer to, amended only by that authority, and enforced by something other than their own goodwill. This paper is about that difference, and about what it takes to make AI governance constitutional rather than merely written down.

It is the most direct argument I make because the distinction it draws is, in my view, the one that separates governance that will hold up when tested from governance that will quietly give way at exactly the moment it is needed.

M Maruf Hossain

SECTION ONE

Executive Summary.

Governance that the governed can override is not governance. This whitepaper makes that argument in full and draws out what it means for how AI authority must be constrained. It rests on three propositions.

First, a rule that yields to the party it governs is not a constraint on that party. This is definitional, not rhetorical. If the people operating AI can set aside its limits whenever a limit becomes inconvenient — and do so without returning to a higher authority or leaving a record — then those limits do not govern them. They describe a preference that holds until pressure arrives. Most AI governance frameworks are built this way: as policies adopted by management, amendable by management, and therefore binding on no one with the power to change them.

Second, a constitution differs from a policy in kind, not in degree. The difference is not that a constitution is stricter or better written. It is that the rules live somewhere the governed cannot reach: adopted by the board rather than management, amendable only by formal resolution rather than at operational discretion, and binding on the very people who operate the systems it governs. A policy is an instrument of management. A constitution is an instrument above management. The distinction is categorical and determines whether a limit holds when tested.

Third, a constitution requires separation of powers and structural enforcement. Two properties make the difference real rather than nominal. The first is separation of powers: those who operate AI cannot also be those who set its limits, because no party can meaningfully constrain itself. The board sets the ceiling; management operates beneath it; neither role collapses into the other. The second is structural enforcement: the limit must be enforced by the infrastructure that runs the AI, not by the goodwill of the people running it — because a limit that depends on goodwill is, once again, only a preference.

This paper develops each proposition, contrasts the policy and the constitution as categories of instrument, and sets out the two structural properties — separated powers and enforced limits — that distinguish governance that holds from governance that merely exists on paper. It closes with five questions that test which kind of organisation it actually has.

SECTION TWO

What Makes *Governance Govern.*

Begin with a question that sounds abstract but is entirely practical: what makes a rule a rule? Not what makes it wise, or well-drafted, or widely endorsed — what makes it actually govern the conduct it addresses, rather than merely describe how that conduct ought to go.

The answer is that a rule governs a party only if that party cannot set it aside at will. A limit you can lift whenever it constrains you is not a limit on you; it is a note about your intentions. This is true of any rule, in any domain. A speed limit that drivers could raise whenever they were in a hurry would not be a speed limit. A budget that a manager could increase whenever they wished to spend more would not be a budget. The constraint exists only to the extent that the constrained party *cannot* remove it unilaterally.

Applied to AI governance, this is decisive. Most frameworks place limits on AI in a policy owned by the same management function that operates the AI. When a limit in that policy becomes inconvenient, management can revise the policy — and does, often for entirely defensible reasons. But the moment the constrained party has the power to relax the constraint, the constraint ceases to govern them. It has become a record of current intention, revisable as intention changes. The framework may be detailed, principled, and sincerely meant. It still does not govern, because it cannot withstand the party it is meant to bind.

This is why the question is not whether an organisation has AI governance, but whether what it has can actually refuse. A governance arrangement that has never once prevented a deployment the business wanted, never once held a limit against commercial pressure, and could not do so even in principle because the people under pressure are the people who hold the pen — that arrangement has not yet demonstrated that it governs anything. It has demonstrated only that it documents.

A limit you can lift whenever it constrains you is not a limit on you. It is a note about your intentions.

SECTION THREE

The Policy *That Yields.*

To see why the ownership of a rule matters so much, follow what happens to a policy at the moment it is tested. The failure is not dramatic. It is quiet, reasonable, and almost invisible — which is precisely what makes it dangerous.

A regulated enterprise has a policy stating that AI systems with a certain level of authority require specific approval before expanding in scope. A competitor launches. The business wants to match the move quickly, and the approval would take time the market will not give. Someone correctly observes that the policy is internal — the organisation wrote it, so the organisation can adjust it. The threshold is raised, or an exception is granted, or the requirement is reinterpreted. The launch proceeds. No rule was broken because the rule was changed first. Nothing in the framework recorded that a constraint had just given way under pressure, since it treated the change as ordinary policy maintenance.

Every step in that sequence is individually reasonable, and that is the point. The policy did not fail because anyone behaved badly. It failed because it *could* yield, and so, under pressure, it did. A policy owned by management is responsive to management's needs by design — that responsiveness is a virtue in an operational policy and a fatal flaw in a governing constraint. The same flexibility that makes a policy easy to keep current makes it unable to hold when holding is the entire point.

The deeper problem is that the yielding leaves no trace. Because amending a policy is a routine act, there is no record that what occurred was a constraint of a constraint under commercial pressure rather than a routine update. A year later, no one can easily reconstruct that the organisation's risk posture was changed in a hurry to chase a competitor. The absence of a constraint and the quiet removal of one look identical in the record. Governance that can be set aside without a trace cannot even tell you, afterwards, that it was set aside.

This is the failure mode that defines the category. A policy that yields is not a defective constitution; it is not a constitution at all. It is operating exactly as a policy operates. The error is not in the drafting. It is in expecting an instrument owned by the governed to constrain the governed.

SECTION FOUR

What
‘Constitutional’ Means.

If a policy cannot govern those who own it, the remedy is not a better policy. It is an instrument of a different category — one whose defining properties are precisely the ones a policy lacks. That category is the constitutional one, and it is worth being exact about what the word means here.

To call a governance instrument constitutional is to claim three things about it. It is *adopted from above*: established not by the management it governs but by the board to which management answers. It is *entrenched*: amendable only through a defined, deliberate procedure — a formal board resolution — rather than at operational discretion, so that changing it is itself a governed act that leaves a record. And it *binds the governed*: it applies to and constrains the very people who operate the systems, including when they would prefer otherwise. A policy may have none of these properties. A constitution must have all three.




Figure 1 — Policy and constitution are different categories of instruments, distinguished by who adopts them, how they are amended, and whether they bind those who operate the system.

The figure is not a comparison of a weak instrument with a strong one. It is a comparison of two *kinds* of instruments. The columns differ on questions of authorship and authority, not on questions of rigour. One could write an extraordinarily detailed and demanding policy, and it would remain, categorically, a policy — revisable by its owners, silent when changed, binding only on those below the authorship. And one could write a brief constitutional constraint that governed far more effectively, because its force comes from where it sits, not from how much it says.

This is the constitutional difference in its plainest form. The strength of a governing limit is not a function of its wording. It is a function of whether the people it constrains can remove it. Move the limit out of their reach — into a board-adopted, formally entrenched instrument — and a short sentence governs. Leave it within their reach, and no amount of detail will make a policy hold.

SECTION FIVE

Separation of Powers.

The first of the two properties that make a constitution real is separation of powers — the principle that the party which operates a capability cannot also be the party that sets the limits on it. It is the oldest idea in constitutional design, and it applies to machine authority with unusual directness.

The reasoning is that no party can meaningfully constrain itself. If the function that deploys and runs AI systems also holds the authority to set and revise the limits on those systems, then the limits are self-imposed — and a self-imposed limit is, by definition, removable by the party that imposed it, the moment it becomes inconvenient. This is not a comment on anyone's integrity. The most disciplined management in the world cannot be constrained by a limit it is free to lift; the freedom to lift it is the absence of the constraint, regardless of whether the freedom is ever exercised.

Separation of Powers: The Operator Cannot Set Its Own Limit



Figure 2 — The board sets the ceiling; management operates beneath it; the infrastructure enforces it. No single party operates the AI and holds the power to raise its own limit.

The separation is therefore between three roles. The board sets the ceiling on autonomy and is the only party that can raise it. Management operates AI systems beneath that ceiling but cannot move it. And the infrastructure enforces the ceiling in operation, so that the limit is not merely declared but applied. The essential property is the one annotated in the figure: authority flows downward, but the power to raise the ceiling flows only from the board. The operator draws against a limit it cannot itself adjust.

This is what converts a ceiling from an aspiration into a constraint. When the board sets the limit and management cannot alter it, the limit governs management — genuinely, in the definitional sense established earlier — because the constrained party has lost the power to remove the constraint. The separation does not assume management will behave badly. It simply ensures that the limit does not depend on management choosing, under pressure, to behave well.

SECTION SIX

The Board-Adopted *Ceiling.*

Separation of powers requires something concrete for the board to actually hold and management to actually operate beneath. In the governance of AI, that concrete object is a ceiling on the total autonomous authority the organisation delegates to its machines — adopted by the board, and amendable only by the board.

The mechanism that makes this workable is a portfolio-level autonomy budget (Hossain, 2026a). Rather than each AI system's authority being approved in isolation by the function that wants to deploy it, the decision authority across the estate is quantified, aggregated into an enterprise total, and capped by a ceiling the board adopts. Individual deployments draw against that ceiling. The board does not need to involve itself in every operational decision; it needs only to hold the one constitutional lever that matters — the total and exclusive power to raise it.

This is what gives the constitutional arrangement its teeth without making the board a bottleneck. Management retains full operational freedom beneath the ceiling: it can deploy, adjust, and run systems as the business requires, drawing against the authorised total. What it cannot do is exceed the total or quietly raise it. To cross the ceiling requires returning to the board for a formal resolution — a deliberate, recorded act of amendment, which is exactly the entrenchment that distinguishes a constitution from a policy. The limit is not administered by those it binds.

The contrast with the policy that yields is now exact. Under a policy, the threshold that stood in the way of a hurried expansion could be revised by those who wanted to expand, silently and as routine maintenance. Under a board-adopted ceiling, the same expansion either fits within the authorised total or it does not — and if it does not, it cannot proceed without a board resolution recording the decision to assume more autonomous authority. The pressure that would have quietly bent a policy now meets a limit that can only be moved in the open, by an authority the business cannot overrule.

SECTION SEVEN

Enforcement by Structure, *Not Goodwill.*

Separation of powers and a board-adopted ceiling establish where the limit lies and who may change it. The second property that makes a constitution real concerns how the limit is applied in the moment: it must be enforced by the structure that operates the AI, not by the goodwill of those operating it.

The reason returns to the argument's first principle. A limit that depends on people choosing to honour it is a limit they can choose not to honour — which is to say, not a limit at all but a preference, however sincerely held. If the board adopts a ceiling but the only thing standing between the organisation and an overshoot is that staff remember to check and elect to comply, then the ceiling has been quietly returned to the discretion of the governed. Constitutional force requires that compliance not be optional at the point of action.

Structural enforcement means the limit is built into the infrastructure that runs the AI, so that an action exceeding the authorised boundary cannot proceed — not because someone declines to permit it, but because the system is constructed so that it does not occur. Two of the accompanying preprints describe mechanisms of this kind: a write-before-execute standard (Hossain, 2026c), under which the authorising record must be committed before an action is permitted, so that an unauthorised action cannot execute; and continuous admissibility (Hossain, 2026b), under which a delegation whose conditions have drifted outside its authorised profile is flagged and halted rather than left to run on goodwill. In each case, the constraint is enforced by construction, not by compliance.

This is the property that closes the gap between a constitution on paper and a constitution in force. A board may adopt the most carefully separated, formally entrenched ceiling imaginable, but if that ceiling is enforced only by people remembering to respect it, it has the structure of a constitution and the force of a policy. Enforcement by structure is what makes the limit binding at the moment of action — which is the only moment at which a limit on machine authority actually matters.

Taken together, the three elements compose the constitutional difference. The limit is adopted from above and entrenched, so the governed cannot rewrite it. Powers are separated, so the operator cannot raise its own ceiling. And enforcement is structural, so compliance is not left to goodwill. An arrangement with all three governs. An arrangement missing any one of them, however elaborate, documents.

SECTION EIGHT

Five Questions *for Your Board.*

The test of whether an organisation has constitutional governance or merely a governance document is not the document's comprehensiveness. It is whether the arrangement could refuse the business when refusing is warranted. Each of the following five questions probes that, and each should be answerable for the AI governance an organisation has in place today.

1. **Can it refuse?** Has your AI governance ever prevented something the business wanted — or could it, even in principle, given who holds the power to change it?
2. **Who owns the limits?** Are the limits on AI authority adopted by the board, or by the same management function that operates the AI and can revise them at will?
3. **Can the operator raise its own ceiling?** Can the function that deploys AI also expand its authority without returning to the board — and if so, what is the limit actually constraining?
4. **Does changing a limit leave a trace?** If a constraint were relaxed under commercial pressure, would the record show that a limit had been moved — or would it look identical to routine maintenance?
5. **Is the limit enforced by structure or by goodwill?** At the moment of action, is an over-limit step prevented by the infrastructure — or only by people remembering to check and choosing to comply?

Governance that can be set aside without trace cannot even tell you, afterward, that it was set aside.

ABOUT

42 Consulting.AI.

42 Consulting.AI builds constitutional AI governance for regulated enterprises. Its core work, the MANDATE Suite, is a comprehensive governance architecture comprising 64 normative instruments across five frameworks — governing the delegation of machine decision authority, the ethical and rights boundaries on its use, the management system that maintains it, and the regulatory extensions for EU and group structures. The Suite is designed for organisations in financial services, insurance, superannuation, healthcare, government, and other regulated sectors operating autonomous AI systems at scale.

The architecture is supported by an academic preprint series published on Zenodo, which establishes the theoretical basis for each governance mechanism, and by the forthcoming book *Governance-First AI*. The Suite's distinguishing conviction is that AI governance at the depth regulated enterprises require cannot be delivered by a compliance framework alone — it requires a constitutional architecture, adopted by the board, enforced by the infrastructure, and continuously tested against current conditions.

Dr M Maruf Hossain, PhD, GAICD · Chief AI Strategist

enquire@42consulting.ai · www.42consulting.ai

The preprint series is available at: zenodo.org/communities/mandate/records

© 2026 42 Consulting.AI · All rights reserved · This whitepaper is general commentary and is not legal, financial, or compliance advice. Organisations should obtain their own professional advice on the application of the referenced governance mechanisms.

REFERENCES

References *and Sources.*

The governance mechanisms described in this paper are set out in full in the MANDATE preprint series, published through Zenodo. The all-versions digital object identifiers below resolve to the current version of each paper.

MANDATE PREPRINT SERIES

- Hossain, M. M. (2026a). *The Autonomy Budget: A Portfolio-Level Framework for Governing Delegated Machine Authority in Regulated Enterprises*. Zenodo. <https://doi.org/10.5281/zenodo.20480491>
- Hossain, M. M. (2026b). *Continuous Admissibility: A Temporal Governance Framework for Delegated Machine Authority in Regulated Enterprises*. Zenodo. <https://doi.org/10.5281/zenodo.20580716>
- Hossain, M. M. (2026c). *Write-Before-Execute: A Governance Standard for Evidentially Defensible AI Decision Logging in Regulated Enterprises*. Zenodo. <https://doi.org/10.5281/zenodo.20603712>



Consulting.AI

Scaling Intelligence

enquire@42consulting.ai · www.42consulting.ai