A General Theory of Recursive Governance

Dependency, Sovereignty, and the Governance of Emerging Technologies

Brian Couzens
SITG Consulting
brian.couzens@sitg-consulting.com

Abstract

Governance of emerging technologies is conventionally organised around individual capability domains. In quantum technologies, attention concentrates on cryptographic disruption and Post-Quantum Cryptography migration. In artificial intelligence, on model safety and alignment. In autonomous systems, on operational control. This paper argues that beneath these domain-specific concerns lies a single structural problem that recurs across every emerging technology: organisations progressively become dependent upon capabilities, measurements, models, and assertions that they can no longer independently reproduce, validate, or meaningfully challenge.

When this condition forms, the primary governance challenge is no longer the technology. It is the maintenance of effective governance when critical decisions depend upon systems the institution cannot independently observe, verify, or control.

The paper develops a general theory of recursive governance built on a single foundational progression: capability becomes dependency, dependency creates trust requirements, trust requirements generate governance obligations, and governance obligations expose sovereignty gaps that concentration amplifies. From this progression, five constructs are derived: Recursive Trust, Recursive Assurance, Measurement Sovereignty, Governance Sovereignty, and Recursive Risk. These are integrated into the Quantum Cyber Risk Framework (QCRF), a recursive governance architecture, together with an accompanying capability model and assessment scaffold.

Two theses are advanced. The *Recursive Governance Thesis* holds that risk emerges through recursive interactions between observation, measurement, validation, assurance, trust, governance, and adaptation, and therefore cannot be understood through analysis of isolated technological, operational, or organisational conditions. The *Sovereignty Thesis* holds that governance resilience is constrained by the degree of independent authority an institution retains over the recursive mechanisms that produce its evidence, confidence, trust, decisions, and adaptation. The framework also addresses recovery: because loss of sovereignty corrupts the faculties an institution would use to detect and repair it, recovery from full governance dependency cannot be generated from within, and requires an exogenous shock, an external authority, or a sovereignty faculty deliberately retained. The central claim is given a single, time-bound, falsifiable predictive test.

The theory is positioned explicitly within the cybernetics and systems-governance tradition. Its contribution is not the discovery of governance feedback, which is well established, but the linkage of sovereignty-reducing dependency to sovereignty loss: the argument that the governing system's authority over its own recursive mechanisms, which that tradition largely brackets as a stable background condition, is in fact progressively eroded by the capabilities the institution adopts. Quantum technologies serve as the proving ground because they expose this erosion earlier and more clearly than any prior domain. The theoretical structure is proposed as general; its generality across domains beyond quantum and artificial intelligence is a claim this paper advances and that requires empirical testing, not one it treats as already established.

Keywords: recursive governance, governance sovereignty, measurement sovereignty, dependency formation, emerging technology governance, systemic risk, concentration, cybernetics, assurance, trust

Contents

1 Introduction: The Problem Beneath the Domains

Each wave of emerging technology arrives with its own governance discourse. Quantum technologies are discussed through the lens of cryptographic disruption and the migration to Post-Quantum Cryptography [11, 12]. Artificial intelligence is discussed through model safety, alignment, and explainability. Autonomous systems are discussed through operational control and liability. Each discourse is organised around the distinctive technical properties of its domain, and each treats the governance problem as a function of those properties.

This paper proposes that the domain-specific framing obscures a structural problem common to all of them.

The common problem is not what these technologies can do. It is what happens to governance when organisations become dependent upon them. Successful capabilities do not remain isolated tools. They become embedded within services, supply chains, decision processes, regulatory frameworks, and infrastructure. Over time, the institutions that adopt them become reliant on outputs, measurements, optimisation decisions, and assertions that they cannot independently reproduce or verify. The capability becomes infrastructure, and the institution's authority over its own governance mechanisms quietly erodes.

The central claim of this paper can be stated directly.

> *The significance of a capability should be measured not only by what it can achieve, but by the dependencies it creates, the trust relationships it introduces, the governance structures required to support it, and the concentration pressures that emerge as adoption increases.*

This claim applies to quantum technologies. It applies equally to artificial intelligence, autonomous systems, advanced analytics, and capability domains not yet mature. It is a claim about the structure of technological dependency, not about any specific technology.

When organisations lose the ability to independently measure, validate, assure, trust, and govern the systems upon which they increasingly depend, the governance problem becomes general. That loss, not the underlying capability, is the subject of this paper.

To prevent the ambiguity that troubles much writing in this area, four roles are fixed at the outset and held consistently throughout. **Emerging-technology dependency** is the problem. **Recursive governance** is the theory advanced to explain it. The **Quantum Cyber Risk Framework (QCRF)** is the framework through which the theory is operationalised. **Quantum technology** is the proving ground: the domain in which the problem appears earliest and most clearly, and through which the framework was developed. The paper is not a quantum risk paper, not a cryptography paper, and not a compliance methodology. It is a governance theory, with quantum as its sharpest case. The title claims generality; the evidence is grounded primarily in quantum and artificial intelligence, with autonomous systems, synthetic biology, and other domains invoked structurally rather than empirically. That gap between the ambition of the claim and the breadth of the evidence is acknowledged here at the outset, and is treated in detail in Section 13, so that no reviewer need raise it before the paper raises it itself.

A word on the name. The framework is called the Quantum Cyber Risk Framework because it was developed through the analysis of quantum-era governance risk, and the name records that origin. The name should be read historically, not as a restriction of scope. The framework's constructs, the dependency–sovereignty linkage, the sovereignty boundary, and the recursive cycle, are domain-agnostic, and the framework applies wherever sovereignty-reducing dependency forms, quantum being the first and clearest domain in which it does. Where the distinction

matters, the reader may take QCRF to denote the general framework and understand its quantum application as the founding instance from which the general form was abstracted.

The argument proceeds from the general to the specific and back. Section 2 establishes the dependency progression that is the foundation of the theory. Section 3 situates the work within the cybernetics and systems-governance tradition and states precisely what it adds. Sections 4 through 7 derive the five constructs from the dependency progression. Section 8 integrates them into the QCRF recursive governance architecture. Sections 9 and 10 present the capability and assessment models. Section 11 develops quantum as the proving ground. Section 12 examines theoretical implications. Section 13 states boundaries and limitations. Section 14 advances the general theory.

2 The Dependency Progression

Technological capability alone rarely creates systemic risk. A capability that is developed but not adopted creates no exposure. Risk emerges when capability becomes embedded within operational environments, decision systems, infrastructure, and economic activity.

The progression is structural, and it is the backbone of the entire framework. It is shown in Figure 1, which also indicates where each construct developed later in the paper attaches to the chain. The figure is referred back to throughout: each construct exists to explain one transition in this progression.

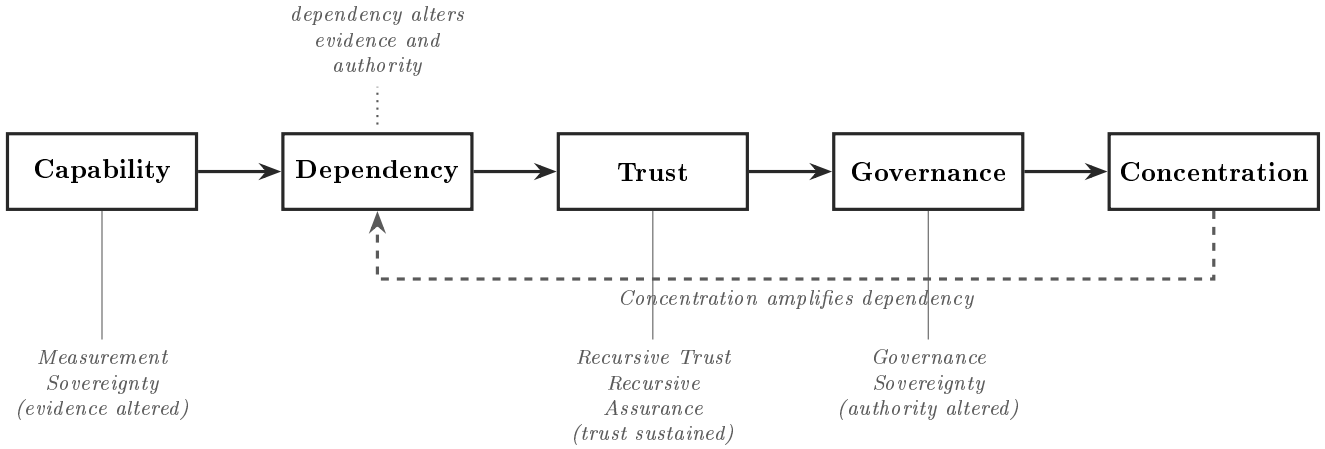

Figure 1: The dependency progression and its derived constructs. Capability becomes dependency; dependency creates trust requirements; trust requirements generate governance obligations; concentration amplifies the chain. Each construct developed in this paper attaches to a transition: Measurement Sovereignty to the point at which dependency alters the institution's evidence; Recursive Trust and Recursive Assurance to the sustaining of trust; Governance Sovereignty to the point at which dependency alters the institution's authority. The progression is the foundation; the constructs are its consequences.

The mechanism runs as follows. Capability creates value. Value encourages adoption. Adoption creates dependency. Dependency creates trust requirements, because a dependent party must rely on something it does not control. Trust requirements generate governance obligations, because reliance that matters must be governed. Governance obligations introduce systemic risk when they cannot be independently discharged, that is, when the institution cannot verify the

basis of the trust it is obliged to govern. Concentration then amplifies the entire chain, because dependency tends to accumulate on a small number of providers, capabilities, and mechanisms.

This progression is the foundation from which every later construct in this paper is derived. Recursive Trust, Recursive Assurance, Measurement Sovereignty, and Governance Sovereignty are not free-standing ideas. They are the constructs required to explain what happens at each transition in the progression and why the transitions matter.

Definition 2.1 (Dependency)**.** *Dependency is the condition in which a capability becomes required for an institution's operation, trust formation, governance, or adaptation, such that the institution can no longer readily function without it. Dependency is observable in technical terms but consequential in governance terms: its significance lies in what it does to the institution's capacity to govern, not merely in the operational reliance it creates.*

This definition deliberately accommodates an objection. A cloud outage that halts operations is plainly a technical dependency with technical consequences. The governance claim is not that the dependency is non-technical. It is that the dependency becomes a governance problem at the point where the institution must rely on, and answer for, a system it cannot independently observe, validate, or control. The technical event is the trigger. The governance condition is what determines whether the institution can detect it, interpret it, and respond.

It must be stated plainly that dependency is not in itself a defect. Dependency creates value, efficiency, and access to capability an institution could not build alone. The division of technological labour is the engine of modern productivity, and an institution that refused all dependency would forgo most of what makes it effective. The argument of this paper is not against dependency. It is about an unrecognised consequence of a particular kind of dependency. The problem is not that institutions depend on others. It is that some dependencies quietly transfer authority over the institution's own governance mechanisms, and that this transfer typically goes unmeasured.

A corollary deserves emphasis, because its neglect has caused persistent misreadings of the sovereignty constructs that follow. Outsourcing a function, engaging a consultant, or relying on a specialist provider does not, in itself, reduce sovereignty. What reduces sovereignty is the loss of the capacity to independently verify, contest, and if necessary replace what is relied upon. An institution that outsources its measurement to a third party but retains the instrumentation, expertise, and contractual authority to reproduce, challenge, and substitute that measurement has not lost Measurement Sovereignty. An institution that performs the same function in-house but lacks the competence or independence to challenge its own outputs may have lost it entirely. The locus of the function is not the issue. The locus of the verification capacity is.

This requires a distinction that the rest of the paper relies upon. Not all dependency reduces sovereignty.

Definition 2.2 (Sovereignty-Reducing Dependency)**.** *A dependency is sovereignty-reducing when it transfers, to a party outside the institution's control, authority over one or more of the recursive mechanisms through which the institution generates evidence, establishes confidence, forms trust, or adapts its decisions. A dependency that leaves these mechanisms under independent institutional authority is sovereignty-preserving, however operationally significant it may be.*

The distinction is sharp and testable. A dependency on a commodity input that is transparently specified, independently verifiable, and available from many substitutable sources is operationally significant but sovereignty-preserving: the institution retains the capacity to measure, validate, and challenge what it receives. A dependency on a proprietary, opaque, single-source capability whose outputs the institution cannot independently reproduce or verify is sovereignty-reducing, because the authority to determine what counts as correct has passed outside the institution's boundary. The same magnitude of operational reliance can fall on either side of the line. What

determines the side is not how much the institution depends, but whether the dependency leaves the institution able to govern the mechanism it now relies upon.

The conditions that push a dependency toward the sovereignty-reducing side are: opacity, where the mechanism cannot be inspected; non-reproducibility, where its outputs cannot be independently generated; non-substitutability, where no alternative source exists; and unchallengeability, where the institution cannot contest the mechanism's determinations. Emerging technologies tend to exhibit all four at once, which is why they reduce sovereignty so readily. But the property is a property of the dependency relationship, not of the technology, and the theory that follows concerns sovereignty-reducing dependency specifically. Where the paper refers to dependency without qualification in the context of governance risk, the sovereignty-reducing kind is meant. Figure 2 presents the four conditions as a gradient: the more conditions hold, the further a dependency moves from sovereignty-preserving to sovereignty-reducing.

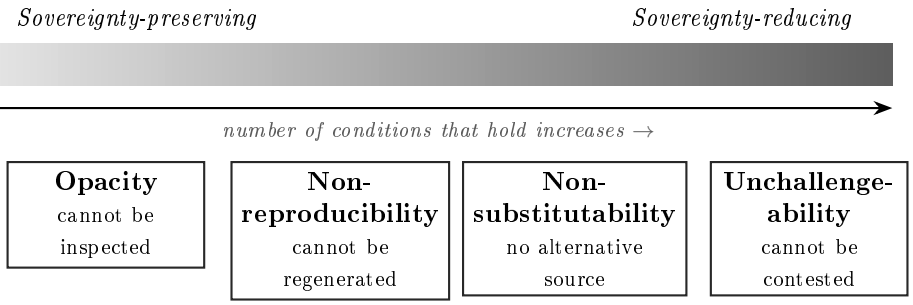

Figure 2: The four conditions of sovereignty-reducing dependency. A dependency exhibiting none of these conditions is sovereignty-preserving however large it is: the institution can still inspect, reproduce, replace, and contest what it relies upon. As conditions accumulate, the dependency moves rightward along the gradient and the institution's authority over the relevant mechanism migrates outside its boundary. Emerging technologies typically present all four at once.

Definition 2.3 (Concentration)**.** *Concentration is the accumulation of dependency upon a small number of actors, capabilities, mechanisms, or infrastructure components. Concentration is simultaneously an economic phenomenon and a governance phenomenon. As a market condition it is well described by existing economic theory. As a governance condition it has a distinct effect: it amplifies dependency, trust asymmetry, governance asymmetry, and sovereignty loss, and it converts localised failures into systemic ones.*

This definition also accommodates an objection. Concentration is, of course, a market phenomenon, and the economics of market power are not in dispute here. The point is that concentration has governance consequences that market analysis does not capture: when trust, measurement, and assurance concentrate on a few actors, the institution's recursive governance cycle becomes dependent on those actors in ways that survive even a perfectly competitive account of the market.

Why does dependency concentrate? Two mechanisms operate, and they reinforce each other.

The first is the verification gap. The same properties that make a dependency sovereignty-reducing, opacity, non-reproducibility, non-substitutability, and unchallengeability, also raise the barrier to entry for alternative providers. A capability that cannot be independently reproduced cannot be easily replicated by a competitor, so supply concentrates. A mechanism whose outputs cannot be independently verified accrues trust to the incumbent that holds the specialised expertise, so trust concentrates. Concentration in emerging-technology markets is therefore not merely the ordinary tendency of returns to scale; it is partly produced by the verification gap itself.

The second is operational friction. In large, distributed operational environments, a capability that has been embedded across multiple systems, regions, processes, and contractual layers becomes structurally irremovable regardless of market conditions. Replacing a vendor whose outputs are woven three layers deep into an institution's operational stack requires a multi-system, multi-region overhaul that no operational team has the capacity or the mandate to execute under normal conditions. The dependency deepens not because no alternative exists in the market, but because the cost of extraction exceeds what the institution can absorb while continuing to operate. This friction is as powerful a concentration driver as any market force, and it operates independently of the verification gap: even a transparent, reproducible capability becomes sovereignty-reducing if its operational embedding makes it practically irremovable.

The two mechanisms compound. The verification gap concentrates supply and trust on a small number of providers. Operational friction locks the institution into whichever provider it adopted first. Together they produce a cycle that is self-reinforcing: dependency drives concentration, and concentration deepens dependency by eliminating both the market alternatives and the operational capacity through which an institution might recover sovereignty. What breaks the cycle is the restoration of independent verification capacity combined with deliberate architectural limits on embedding depth, which is precisely what the sovereignty constructs of this paper are intended to make visible and assessable.

Definition 2.4 (Governance Transformation)**.** *Governance Transformation is the change in the character of an institution's governance from substantive control to formal dependence, occurring as sovereignty-reducing dependency accumulates. It is not the failure of governance activity, which may continue undiminished in volume, but a change in its nature: decisions are still made, reports still produced, and assurances still issued, while the authority to determine the evidence, validation, and trust on which those activities rest has migrated outside the institution. Governance Transformation is the dependent variable that the dependency–sovereignty linkage seeks to explain.*

Proposition 2.1 (Dependency Drives Governance Transformation)**.** *The rate at which governance is transformed is determined by the rate of sovereignty-reducing dependency formation, not by the rate of capability development. A capability that develops rapidly but is adopted in a sovereignty-preserving manner transforms governance slowly. A capability adopted in a sovereignty-reducing manner transforms governance at the rate of that adoption, regardless of how mature the underlying technology is. Governance must therefore track sovereignty-reducing dependency formation, not technological progress.*

Governance Transformation is best understood as a progression of states rather than a single event, shown in Figure 3. The danger of the progression is that it is largely invisible from inside: the volume and form of governance activity are preserved at every stage, so an institution can pass through all four states while its dashboards, committees, and attestations appear unchanged.

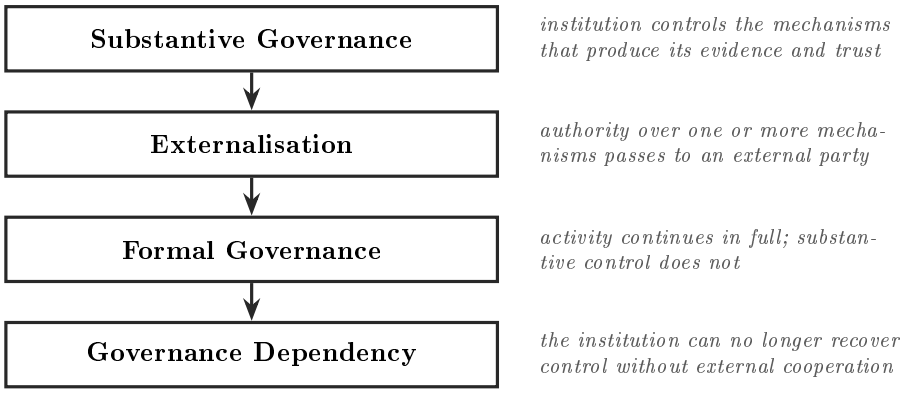

Figure 3: The Governance Transformation progression. Substantive governance degrades to formal governance not through the cessation of activity but through the externalisation of authority over the mechanisms on which that activity depends. By the final state, governance dependency, the institution retains the appearance and routines of control while having lost the capacity to recover it unaided. Because the visible form of governance is preserved throughout, the transformation is rarely detected from within until a failure exposes it.

3 Intellectual Lineage and Contribution

The claim that governance involves feedback is not new, and this paper does not present it as new. Recursive and self-correcting governance has a deep intellectual history, and the contribution of this work can only be stated honestly by placing it against that history.

3.1 What the Tradition Established

Cybernetics established the foundations. Ashby's principle of requisite variety holds that a regulator must possess at least as much variety as the system it regulates, or it cannot control it [1]. This is a direct ancestor of the present argument: a governing institution that cannot match the complexity of the system it governs will fail to govern it. Beer's Viable System Model formalised the idea of nested, self-referential organisational layers, each capable of regulating itself while contributing to the viability of the whole [2]. The recursive structure of governance, in which each level reproduces the regulatory pattern of the level above, is Beer's, not this paper's.

Control engineering and management contributed the cycle. The Plan-Do-Check-Act cycle in quality management, and the Observe-Orient-Decide-Act loop in decision theory, both describe iterative processes in which action produces feedback that informs subsequent action. The QCRF cycle developed in Section 8 shares this iterative form.

Sociology contributed self-reference. Luhmann's theory of social systems describes governance and law as self-referential systems that produce the very elements from which they are composed [3]. The notion that a governing system observes, and is changed by, its own outputs is Luhmann's.

Safety and resilience science contributed failure dynamics. Perrow's analysis of normal accidents showed how tightly coupled, complex systems generate failures that are systemic rather than attributable to single causes [4]. Resilience engineering established that safety is a dynamic property maintained through continuous adaptation, not a static state achieved through controls [5]. Power's analysis of the audit society showed how verification can become ritual, producing the appearance of assurance without its substance [6].

These traditions are not background. They are the foundation. Any honest account of recursive governance must acknowledge that the feedback structure, the nesting, the self-reference, and the failure dynamics were established by others.

3.2 What This Paper Adds

The contribution of this paper is not the feedback loop. It is a claim that the existing tradition does not make.

The cybernetics and systems-governance tradition does not centrally problematise the governing system's authority over its own regulatory mechanisms; it largely brackets the question, treating that authority as a stable background condition. Ashby's regulator has its variety. Beer's system has its operations. The PDCA cycle proceeds on the basis that the organisation controls its own measurement and correction. In each case, the governing system's authority over its own recursive machinery is held constant while the analysis attends to how that authority is best exercised. The question the tradition foregrounds is how to use the authority well. The question it does not foreground is what happens when the authority itself is lost.

This paper asks that second question: what happens when that authority erodes?

The central contribution is the linkage of sovereignty-reducing dependency to sovereignty loss. As an institution becomes dependent, in the sovereignty-reducing sense of Definition 2.2, on external capabilities, it progressively loses authority over the very mechanisms the tradition held constant. It loses authority over measurement, because it can no longer independently observe system state. It loses authority over validation, because it cannot challenge the models that determine correctness. It loses authority over assurance, because it cannot independently verify the evidence of confidence. The regulator's variety is not merely insufficient, in Ashby's sense; it is being externalised. The viable system's recursive autonomy, in Beer's sense, is being transferred to actors outside its boundary.

This is the gap. The tradition explains how recursive governance works when the institution controls its own mechanisms. It does not centrally address what happens to recursive governance when adoption of external capability strips the institution of that control. That is what sovereignty-reducing dependency does, and that is what this paper sets out to explain. Measurement Sovereignty and Governance Sovereignty are the constructs that name the property the tradition held constant and this paper shows to be at risk.

Principle 3.1 (The Sovereignty Gap)**.** *Classical governance theory brackets the governing institution's authority over its own recursive mechanisms, treating it as a stable background condition. Sovereignty-reducing dependency removes that stability. The governance problem of emerging technology is therefore not a failure to apply known feedback principles, but the progressive loss of the sovereignty those principles presuppose.*

This gap can be stated as a single linkage, which is the central original claim of the paper and the thread from which every construct that follows is drawn. It is stated here, once, in the cleanest possible form.

Thesis 3.1 (The Dependency–Sovereignty Linkage)**.** *Sovereignty-reducing dependency drives sovereignty loss, and sovereignty loss drives governance transformation. As an institution becomes dependent, in the sovereignty-reducing sense, on a capability it cannot independently reproduce, verify, or challenge, it loses authority over the recursive mechanisms that produce its evidence, confidence, trust, and decisions; and as that authority is lost, its governance is transformed from substantive control into formal dependence. The chain runs: sovereignty-reducing dependency → sovereignty loss → governance transformation. The qualification matters: ordi-*

nary, sovereignty-preserving dependency, however large, does not drive the chain. Everything else in this framework is an elaboration of one of these three terms or of the links between them.

The remainder of the paper develops this linkage. The constructs of Sections 4 through 7 specify what sovereignty loss consists of at each layer. The architecture of Section 8 shows where in the recursive cycle the loss occurs. The assessment scaffold of Section 10 provides the means to detect it. The two theses of Section 12, the Recursive Governance Thesis and the Sovereignty Thesis, are the generalisations of, respectively, the first and second links in the chain. The linkage is shown in Figure 4, which is the intellectual centre of the paper.

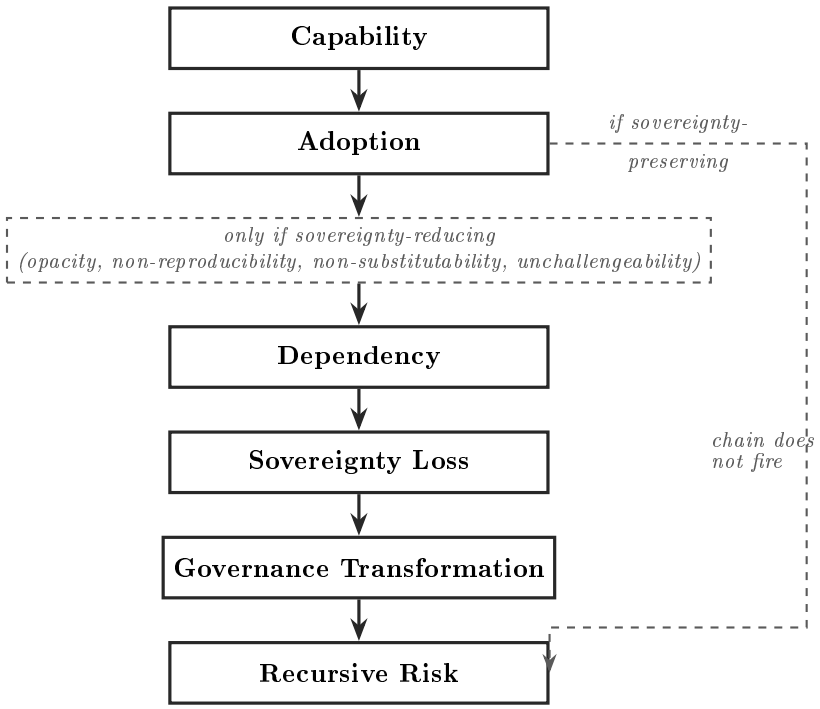

Figure 4: The dependency–sovereignty linkage, the central thesis of the paper. Capability adoption creates dependency, but the chain to governance transformation fires only when the dependency is sovereignty-reducing, that is, when it satisfies the conditions of Figure 2. Sovereignty-preserving dependency, however large, bypasses the chain (dashed path). Where the chain does fire, sovereignty loss transforms governance from substantive control into formal dependence, and the recursive risk that follows compounds through the governance cycle of Figure 5.

4 Recursive Trust

Trust is conventionally treated as a precondition. An institution establishes trust in a system, a provider, or a measurement, and governance proceeds on that basis. Once established, trust is assumed to persist until explicitly revoked.

This framing is inadequate for systems characterised by complexity, interdependence, and continuous change, because the conditions that justified the original trust do not remain fixed.

Definition 4.1 (Recursive Trust)**.** *Recursive Trust is trust that is continuously regenerated, tested, and reconstructed through repeated cycles of measurement, validation, and assurance. It treats trust not as a state to be established once, but as an output of governance that must be reproduced as the conditions that justified it change.*

The distinction changes what governance must do. If trust is a precondition, governance verifies it once and then responds only to explicit violations. If trust is recursive, governance must continuously regenerate the evidence, validation, and assurance on which trust rests. Trust that is inherited from a prior governance state without reconstruction against current evidence is not governance. It is assumption.

Trust becomes fragile under four conditions, each of which is intensified by dependency on external capability. When measurement is externalised, the institution cannot independently observe the state on which its trust claims depend. When validation is opaque, the mechanisms that establish correctness cannot be examined or challenged. When assurance is unchallengeable, the evidence supporting confidence cannot be independently reproduced. When governance is symbolic, the formal apparatus exists without substantive control over the mechanisms that produce trust.

These conditions are not hypothetical for emerging technologies. A quantum sensing output may require equivalent instrumentation to reproduce. A quantum optimisation result may depend on models that cannot be classically verified within useful time. An AI system's outputs may change between assurance cycles without any change to the system, because the model has adapted to new data. In each case the recursive trust cycle is stressed because the evidence on which trust depends becomes harder to independently generate.

Proposition 4.1 (Trust Reconstruction)**.** *Trust cannot be inherited across governance states without reconstruction. As conditions change through capability adoption, dependency formation, or environmental disruption, trust must be regenerated from current evidence, current validation, and current assurance. Inherited trust that is not reconstructed becomes, over time, indistinguishable from unwarranted assumption.*

Because trust is recursive, trust failure propagates. A failure of measurement integrity is not contained as a measurement problem. It undermines the validation that depends on measurement, the assurance that depends on validation, the trust that depends on assurance, and the decisions that depend on trust. This propagation is the mechanism by which a localised technical failure becomes a systemic governance failure.

5 Recursive Assurance

Assurance is conventionally treated as terminal. An audit is conducted, a certification granted, a compliance assessment completed, and the result is treated as settled until the next cycle. This model fails under recursive complexity, because the system being assured changes between cycles and the mechanisms of assurance themselves rest on assumptions that may no longer hold.

Definition 5.1 (Recursive Assurance)**.** *Recursive Assurance is assurance that recognises its own contingency: the condition in which assurance mechanisms are themselves subject to validation, and in which the evidence supporting an assurance claim must be regenerated as conditions change. Assurance, on this view, requires assurance.*

The recursive character becomes visible when an assurance claim is decomposed. The claim depends on evidence. Evidence depends on measurement. Measurement depends on observation. Observation depends on instrumentation. Trust in the instrumentation depends on prior assurance. The loop is not a logical defect; it is the structural reality of assurance in complex systems. It becomes a problem only when it is unrecognised or uncontrolled.

Conventional governance treats the loop as invisible. Certification cycles assume the certified state persists between assessments. Compliance frameworks assume that conformity demonstrated at assessment reflects the operational state between assessments. Audit methodologies

assume that sampling yields representative evidence. Power's account of the audit society describes precisely the risk: assurance that has become ritual, producing the appearance of verification without its substance [6]. Recursive Assurance is the corrective. It insists that the assurance mechanism is itself a system requiring assurance.

Emerging technologies stress each assumption. Quantum systems introduce measurement uncertainty that classical assurance methods were not built to accommodate. Adaptive systems produce outputs that change between cycles without modification to the system. Hybrid architectures create assurance boundaries that span different computational paradigms.

Proposition 5.1 (Assurance Circularity)**.** *Assurance that is not itself subject to independent validation is circular. The practical consequence is that the capacity to independently validate assurance claims, here termed assurance sovereignty, is a prerequisite for governance resilience rather than an optional refinement.*

The point extends beyond technical assurance. Legal, contractual, regulatory, and evidentiary assurance all exhibit the same recursion. A regulatory determination depends on evidence; the evidence depends on measurement; the measurement depends on the regulated system; and the regulated system is influenced by the determination. The loop appears wherever governance interacts with the system it governs.

6 Measurement Sovereignty

Measurement is the foundation of governance. Without independent access to reliable evidence about system state, no governance function can operate with integrity. Policy depends on data. Regulation depends on reporting. Assurance depends on observation. Trust depends on evidence. Every governance decision rests, directly or indirectly, on measurement.

Definition 6.1 (Measurement Sovereignty)**.** *Measurement Sovereignty is the capacity of an institution, sector, or governing authority to maintain independent confidence, control, and assurance over the measurements that influence its critical decisions. It comprises the ability to observe system state, define meaningful metrics, interpret evidence, and challenge measurement claims without dependence on the systems or actors being measured.*

It is worth distinguishing this construct from the managerial commonplace that what gets measured gets managed. That maxim concerns where managerial attention is directed. Measurement Sovereignty concerns who controls the measurement apparatus and what happens when that control is externalised. The two are not the same claim. An institution can measure diligently, direct its attention well, and still lack Measurement Sovereignty entirely, because the metrics are defined by the vendor, the telemetry is controlled by the platform, the interpretation requires the measured system to explain itself, and independent reproduction is infeasible. The managerial maxim assumes the measurement is the institution's own. Measurement Sovereignty asks whether it is.

An institution holds Measurement Sovereignty to the extent that it can independently observe the systems it depends on, define the metrics by which their state is judged, interpret outputs without relying on the measured system for explanation, and challenge measurement claims through independent reproduction or equivalent validation. Loss of Measurement Sovereignty occurs when any of these is externalised.

The consequences propagate through the whole governance system, and they propagate from the base. If measurement cannot be trusted, validation is unreliable; if validation is unreliable,

assurance is unfounded; if assurance is unfounded, trust is misplaced; if trust is misplaced, decisions are unsound; if decisions are unsound, adaptation is maladaptive. Within the architecture presented here, in which measurement sits at the foundation of the recursive chain, its compromise is the most consequential single point of failure, because it corrupts the evidence feeding every higher layer.

The derivation should be stated rather than left to inference, because the claim is load-bearing. The recursive cycle is not a ring of equals. It is a directed chain in which each layer operates on the output of the layer below it: validation validates measurements, assurance assures validations, trust rests on assurances, governance acts on trust, and adaptation modifies the system that will be measured next. Corruption at the base, measurement, propagates upward through every subsequent layer because every subsequent layer takes measurement's output as its input. Corruption at a higher layer, say governance capture or adaptation failure, is catastrophic but contained: it corrupts the layers above it and feeds error into the next cycle, but the layers below it retain their integrity until that error reaches them through the recursive return. Governance capture corrupts decisions and adaptation; it does not corrupt the evidence already gathered. Adaptation failure prevents learning; it does not falsify the observations already made. Measurement failure does something worse: it substitutes the evidence itself, so that every layer above it operates on a false basis without any internal signal that the basis is false. That is why measurement is the foundation: not by assertion, but because the directed structure of the chain means corruption at the base has the widest propagation and the least visibility.

This is a claim about the architecture, not a universal ranking of failure severity. Governance capture and adaptation collapse can be catastrophic by routes that bypass the measurement layer entirely, and this paper does not deny it. The claim is structural: within a recursive governance cycle of this form, the base layer is the single point whose compromise is hardest to detect from within and widest in its effects.

Proposition 6.1 (Measurement as Governance Foundation)**.** *A system cannot be governed with independent authority if it cannot be independently measured. Measurement Sovereignty is therefore not an operational convenience but a precondition of sovereign governance. The loss of Measurement Sovereignty is the precise point at which governance passes from substantive control into formal dependence. Governance of some kind may continue without independent measurement; sovereign governance cannot.*

This proposition is not the claim that measurement is useful, which would be a truism. It is the claim that independent measurement is the dividing line between governing a system with independent authority and merely being answerable for one. It is an inferential claim, grounded in observed patterns of vendor-defined metrics, proprietary validation, and infrastructural dependence, that many institutions which believe they govern their technological dependencies have, by this criterion, already crossed that line. The prevalence of the condition is not established here and is one of the empirical questions the framework raises rather than answers.

A practical qualification is warranted. The proposition establishes independent measurement as a precondition of sovereign governance. It does not follow that every institution must build equivalent measurement capacity for every dependency it holds. The cost and complexity of independent measurement vary by capability, by sector, and by the severity of the governance consequences at stake. An institution whose dependency is operationally significant but whose governance exposure is modest may reasonably conclude that the cost of full independent measurement exceeds the risk it mitigates. The framework's contribution is not to prescribe a particular cost threshold but to make the trade-off visible: an institution that forgoes independent measurement over a decision-critical dependency should understand that it is accepting formal governance over that dependency, not achieving sovereign governance at lower cost. Provided the trade-off is made consciously and documented, it is a legitimate governance decision. The

danger the framework addresses is the far more common case in which the trade-off is never recognised.

7 Governance Sovereignty

Definition 7.1 (Governance Sovereignty)**.** *Governance Sovereignty is the capacity of an institution or governing authority to maintain independent authority over the recursive mechanisms through which evidence is generated, trust is established, assurance is validated, and decisions are adapted. It is the ability to govern not only outcomes, but the processes that produce governance itself.*

Governance Sovereignty is not a static attribute. It is a dynamic property of control over the measurement, validation, assurance, trust, and adaptation cycles that underpin decision-making. It is, in the terms of Section 3, the property the cybernetics tradition assumed and this paper shows to be erodible.

A note on the term is warranted, because "sovereignty" carries established meanings in political theory, state theory, and international relations, where it denotes the supreme authority of a state over its territory and the recognition of that authority by other states. The present usage is distinct and should not be read through that lens. Here, sovereignty denotes an institution's independent authority over its own recursive governance mechanisms, the capacity to measure, validate, assure, trust, and adapt without dependence on the actors being governed. The word is retained, rather than replaced with a softer alternative such as "authority" or "control," because it alone conveys both the dimensions that matter: not merely the operational ability to act, but the standing to determine, for oneself, what counts as evidence, correctness, and warranted reliance. Where the distinction from political sovereignty needs emphasis, the construct may be read as "epistemic and procedural sovereignty over governance," though the shorter term is used throughout.

A second clarification prevents a misreading that would invert the framework's intent. Governance Sovereignty is a property of the institution's verification capacity, not of its organisational perimeter. It does not require, and should not be read to require, that every governance function be performed in-house. An institution that delegates assurance to an external auditor but retains the instrumentation, data access, and expertise to reproduce, challenge, and if necessary replace that auditor's findings retains sovereignty over the assurance layer. An institution that performs assurance internally but whose internal team lacks the independence or competence to challenge the system it assures does not. The sovereignty question at each layer is not "who performs this function?" but "does the institution retain the capacity to independently verify, contest, and replace the output of this function?" An affirmative answer preserves sovereignty regardless of whether the function sits inside or outside the institutional boundary. A negative answer indicates sovereignty loss regardless of where the function is performed.

7.1 The Limits of Conventional Governance

Traditional governance assumes linearity: information flows upward, decisions flow downward, and authority remains external to the system governed. Emerging technologies break this assumption. Governance decisions alter the system; the altered system produces new evidence; the new evidence reshapes governance assumptions. Governance becomes both observer and participant. It becomes recursive.

Linear governance fails here not for want of intent or resource, but structurally. A model that

assumes stable evidence, stable trust, and stable assurance cannot survive conditions in which all three are continuously modified by the systems being governed.

7.2 Sovereignty Dependencies

Governance Sovereignty rests on four subordinate sovereignties, each a necessary condition for the integrity of the whole, and each defined in parallel terms.

Measurement Sovereignty is the capacity to independently observe and interpret system state; without it, governance depends on external definitions of reality. **Assurance Sovereignty** is the capacity to independently validate the claims of correctness on which confidence rests; without it, governance depends on external claims it cannot test. **Trust Sovereignty** is the capacity to determine, on independent evidence, where reliance is warranted; without it, governance depends on assumptions it cannot verify. **Adaptation Sovereignty** is the capacity to alter the institution's own processes in response to new evidence; without it, governance cannot learn, and responds to new conditions with inherited reflexes.

These four are defined here in parallel for completeness. Two of them, Measurement Sovereignty and Governance Sovereignty, are elevated to full sections in this paper, not because they are more important in principle, but because they mark the two ends of the chain: measurement is the foundation on which the others rest, and governance is the composite that the others constitute. Assurance Sovereignty and Trust Sovereignty are the intermediate conditions, fully real but analytically situated between the two ends, and they are treated at the length their intermediate position warrants. A complete taxonomy of sovereignty, including whether each kind can be retained, shared, delegated, federated, or recovered, is identified as future work in Section 13.

Proposition 7.1 (Sovereignty Collapse)**.** *Loss of sovereignty at a single layer introduces governance asymmetry. Loss across multiple layers produces governance dependency. Loss across the majority produces governance collapse. Collapse here does not mean the cessation of governance activity. It means the continuation of governance activity without substantive control over the mechanisms that determine outcomes.*

7.3 Governance Sovereignty Failure

Sovereignty degrades when critical governance functions are externalised: vendor-defined risk metrics accepted as authoritative without independent validation; opaque assurance mechanisms substituted for transparent verification; regulatory compliance used as a proxy for understanding; the inability to challenge governance assumptions; and dependence on external actors for trust formation. The result is formal governance without substantive control: a system that appears governed but is not sovereign. The distinction between formal and sovereign governance is the distinction between compliance and resilience.

Principle 7.1 (Governance Sovereignty)**.** *A governance system is sovereign only to the extent that it controls the recursive mechanisms through which evidence is measured, confidence is established, trust is formed, and decisions are adapted. Governance Sovereignty is a property of recursive authority, not of organisational hierarchy.*

7.4 The Recovery Problem

The architecture raises a problem that the theory must address directly, because it follows from the theory's own structure. If Measurement Sovereignty sits at the foundation of the recursive

cycle, and its loss corrupts every layer above it, then an institution that has reached the state of governance dependency is in an apparent trap. Its measurement is externalised, so its evidence is suspect; its validation rests on that evidence; its assurance rests on that validation; its adaptation rests on all three. The faculties it would use to diagnose its own condition and to repair it are precisely the faculties that have been corrupted. A system running on externalised feedback cannot, it seems, use that feedback to detect that it should stop trusting it. This is a bootstrap problem: recovery appears to require the very sovereignty whose absence defines the condition.

The theory's answer is that recovery from full governance dependency is not endogenous. It cannot be bootstrapped from within the corrupted cycle, and the framework does not pretend otherwise. Recovery requires an input from outside the corrupted loop. Three such inputs are possible, and they are worth naming because they are the only routes the structure permits.

The first is **exogenous shock**: a failure large enough that it cannot be absorbed or reinterpreted by the externalised mechanisms, and that exposes the gap between formal governance and substantive control. A vendor collapse, a public incident, a regulatory finding that the institution could not itself have produced. Shock is the most common route in practice and the most costly, because by definition it arrives after the failure rather than before.

The second is **external authority**: an actor outside the institution's corrupted cycle, a regulator, an independent auditor with genuine reproduction capacity, or a supervisory body, that possesses the Measurement Sovereignty the institution has lost and can impose it from outside. This is the rationale for treating independent measurement capability, rather than third-party assurance, as a supervisory concern in its own right: the supervisor must hold the sovereignty the supervised institution cannot.

The third, and the only route available before shock or intervention, is **retained sovereignty at a single layer**. If the institution has preserved genuine independent measurement over even one decision-critical domain, that surviving capacity provides an uncorrupted vantage point from which the rest of the cycle can be rebuilt. This is the practical case for never allowing measurement sovereignty to be lost completely, even where dependency is otherwise deep: a single retained measurement faculty is the seed from which recovery can be bootstrapped, whereas its total loss forecloses internal recovery entirely.

Proposition 7.2 (Non-Endogenous Recovery)**.** *Recovery from full governance dependency cannot be generated from within the corrupted recursive cycle. It requires an input from outside that cycle: an exogenous shock, an external authority that holds the sovereignty the institution has lost, or a sovereignty faculty deliberately retained at some layer and never externalised. The corollary is that the cheapest insurance against irreversible governance dependency is the preservation of independent measurement at a minimum of one decision-critical layer, since total measurement loss removes the only endogenous route to recovery.*

7.5 The Limit Case: Categorically Unverifiable Capabilities

The recovery argument rests on an assumption that should be made visible: that measurement sovereignty is recoverable in principle, even if the recovery is expensive, slow, or operationally difficult. For the capabilities discussed so far, this assumption holds. Quantum sensing outputs are hard to reproduce independently, but not categorically impossible given equivalent instrumentation. AI model outputs are opaque, but independent test harnesses and validation datasets exist. The measurement faculty the institution is advised to retain as a seed is, in each of these cases, something that can in principle be retained.

The harder question is whether there exist, or will exist, frontier capabilities for which this assumption fails: capabilities whose outputs categorically cannot be independently measured, not

because the institution lacks the resources or expertise, but because no independent measurement is possible given the nature of the capability itself. If the measurement faculty the institution is supposed to retain is the one thing it cannot hold onto, the framework's recovery logic reaches its structural limit.

The theory's prediction for this case follows directly from its own architecture and requires no additional apparatus. If measurement sovereignty cannot be retained at any layer, then by Proposition 7.2 no endogenous recovery path exists. If the capability is also opaque, non-reproducible, and non-substitutable, then by Thesis 3.1 the institution enters permanent, irrecoverable governance dependency from the point of adoption. Governance activity continues; substantive governance does not return. The condition is not a temporary state pending recovery. It is a terminal state within the framework's logic.

Whether any current capability truly occupies this position is an empirical question the paper cannot settle. The more likely near-term reality is that capabilities are presented as categorically unverifiable when in fact they are merely expensive to verify, and that the distinction between "cannot be measured" and "will not invest in measuring" is frequently elided to the advantage of the provider. The framework's practical contribution in this borderland is to insist on the distinction: an institution should not accept a provider's claim that independent verification is impossible without testing that claim as rigorously as it would test the capability itself.

But the limit case remains. If a capability genuinely forecloses independent measurement, the framework predicts that its adoption is a permanent sovereignty transfer, made under conditions where the institution's future governance of the capability will be formal rather than substantive, with no structural route back. That prediction may prove rare in practice. It is worth stating because it follows from the theory without modification, and because the institution that understands the prediction before adoption is in a different position from the one that discovers it afterward.

8 The QCRF Recursive Governance Architecture

The five constructs derived above integrate into a single architecture. The Quantum Cyber Risk Framework (QCRF) is a recursive governance architecture for analysing how risk is generated, transformed, and amplified within complex socio-technical systems under conditions of dependency and diminishing sovereignty. Its purpose is not to measure static risk, but to evaluate the quality of the recursive governance mechanisms that determine how risk evolves over time.

8.1 The Architecture

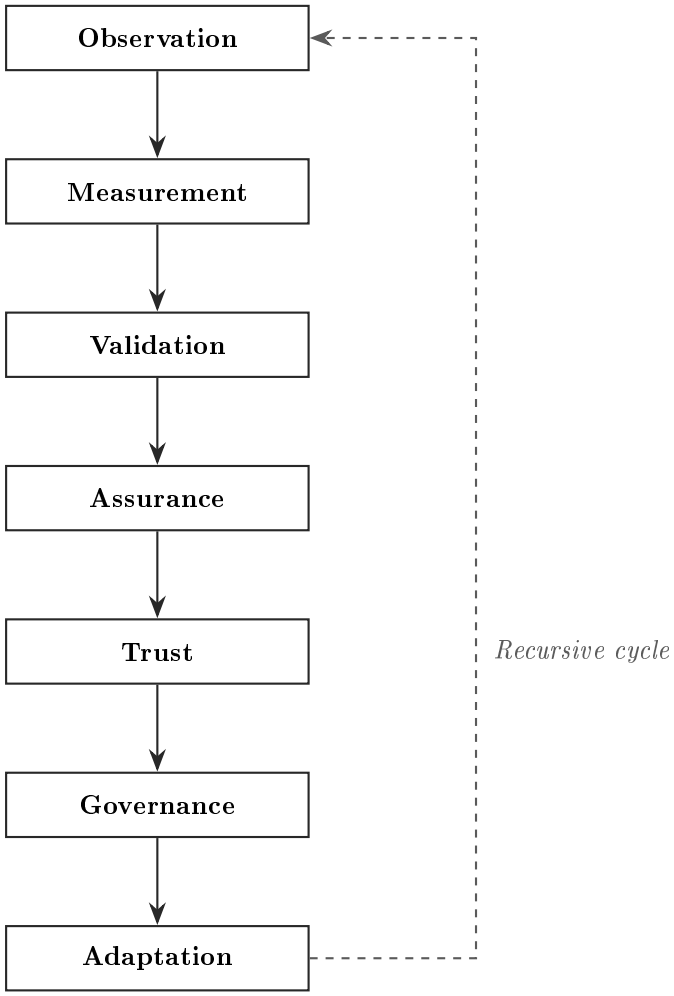

Figure 5: The QCRF recursive governance cycle. Each layer performs a distinct governance function and depends on the output of the layer above. The dashed return path represents the recursion through which adaptation modifies the system, producing the new observations that begin the next cycle. The architecture extends the iterative form common to PDCA, OODA, and the Viable System Model; its distinctive feature is the sovereignty boundary present at each layer, discussed below.

Each layer performs a distinct function. Observation determines what enters the governance system. Measurement determines how evidence is structured. Validation determines how correctness is established. Assurance determines how confidence is formed. Trust determines how reliance is justified. Governance determines how decisions are made. Adaptation determines how governance evolves. The architecture is recursive because each decision modifies the conditions under which future evidence will be observed, measured, validated, assured, and trusted.

8.2 Sovereignty Boundaries

The architecture's distinctive feature, and the point at which it departs from its predecessors, is that each layer contains a sovereignty boundary: a point at which authority over the verification of the function's output may remain internal to the institution or pass outside it. The boundary does not mark who performs the function; it marks who retains the capacity to independently verify

its output. Observation: who can independently verify what is observed? Measurement: who can independently verify how evidence is structured? Validation: who can independently verify correctness? Assurance: who can independently verify confidence? Trust: who can independently verify reliance? Governance: who can independently verify decisions? Adaptation: who can independently verify change?

The classical iterative cycles, PDCA, OODA, and the recursive levels of the Viable System Model, assume these questions are answered internally. QCRF makes the question explicit at every layer precisely because, under dependency, the answer increasingly is not the institution. The sovereignty boundary is where the dependency progression of Section 2 meets the recursive cycle of the cybernetics tradition. Figure 6 maps the boundary across the cycle, showing at each layer the question that decides whether authority remains inside the institution or passes outside it.

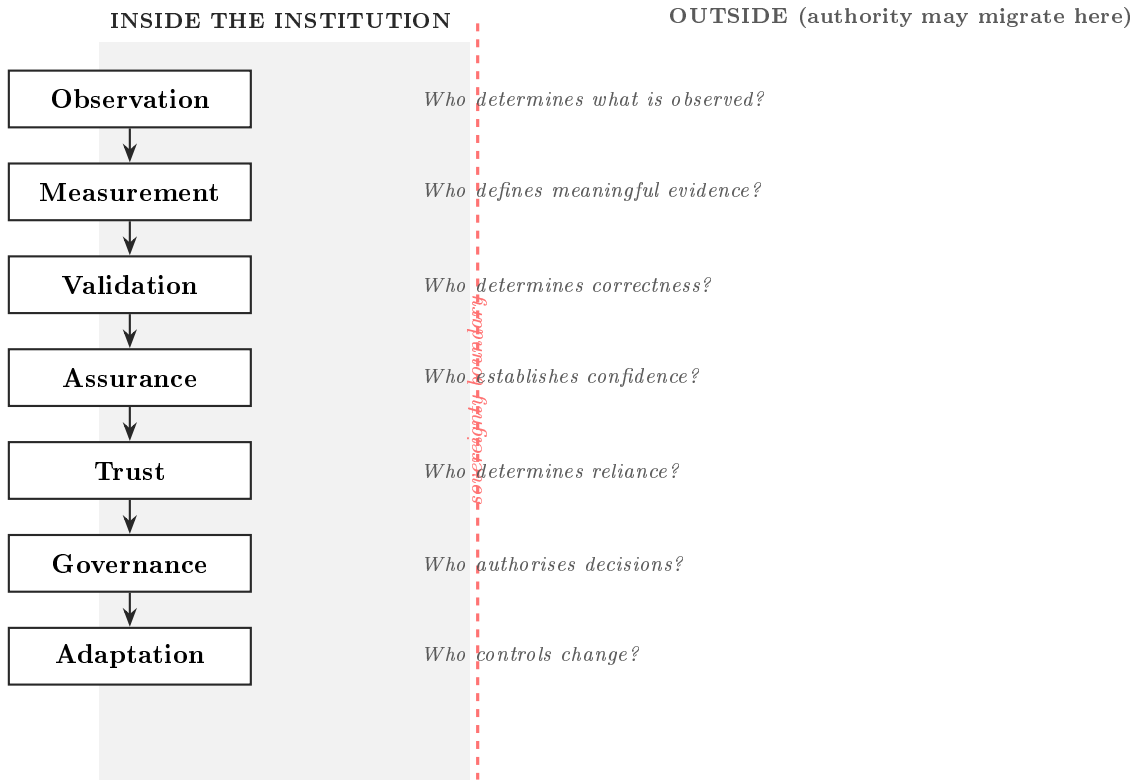

Figure 6: The sovereignty boundary map. At each layer of the recursive cycle, a single question decides on which side of the boundary authority sits. The question is not who performs the function but who retains the capacity to independently verify its output. When the institution holds that capacity, sovereignty over the layer is retained; when an external party holds it exclusively, sovereignty has migrated outside the boundary. Sovereignty-reducing dependency is the mechanism by which the answers shift rightward, one layer at a time, while the cycle continues to run. This boundary, present at every layer, is the feature that distinguishes QCRF from the iterative cycles it extends.

8.3 Risk Propagation

Recursive systems accumulate error. Without sovereign control over the recursion, error compounds through the cycle: observation error becomes measurement distortion, becomes validation failure, becomes assurance weakness, becomes trust misplacement, becomes governance error, becomes maladaptive response, becomes amplified risk in the next cycle. The propagation is not

a metaphor for "garbage in, garbage out." It is a claim about accumulation: because the output of each cycle becomes the input of the next, uncorrected error is not merely passed along but compounded across cycles, and the compounding is fastest precisely where sovereignty over the relevant layer has been lost.

Principle 8.1 (Recursive Risk)**.** *Risk is not a property of events but of recursive interactions. The resilience of a system is proportional to its capacity to maintain sovereign control over the recursive relationships between evidence, assurance, trust, governance, and adaptation under conditions of uncertainty and change.*

Recursive Risk, as stated here, is the risk-focused expression of what Section 12 generalises as the Recursive Governance Thesis. The principle describes the property at the level of the architecture; the thesis states it as a general claim about how risk emerges. They are the same proposition viewed at two levels of generality, and the later thesis should be read as the formal statement of which this principle is the architectural instance.

9 The QCRF Capability Model

The architecture describes how recursive risk propagates. The capability model evaluates the organisational capacities required to manage that propagation. It shifts the focus of evaluation from controls to capabilities, from static compliance to recursive resilience, and from technical artefacts to sovereign governance functions.

9.1 Capability Domains

Seven capability domains correspond to the seven layers of the architecture.

Layer	**Capability**	**Purpose**
Observation	Observability	Detect system state
Measurement	Measurement	Quantify system state
Validation	Validation	Confirm correctness
Assurance	Assurance	Establish confidence
Trust	Trust	Maintain justified reliance
Governance	Governance	Make sovereign decisions
Adaptation	Adaptation	Respond to change

Table 1: QCRF capability domains mapped to architecture layers.

Each domain is evaluated through its purpose, the sovereignty on which it depends, its characteristic failure modes, the way capability disruption amplifies those failures, and the questions an assessor must ask to judge it.

9.2 Characteristic Failure Modes

Observability: blind spots; externalised telemetry; inability to detect state change without vendor instrumentation. *Measurement*: unverifiable or vendor-defined metrics; quantification that cannot be independently reproduced. *Validation*: unchallengeable logic; opaque or proprietary models; validation that depends on the system being validated. *Assurance*: claims that cannot be independently tested; circular assurance in which evidence is generated by the assured system. *Trust*: assumptions not reconstructed against current evidence; trust asymmetry in which

the trusted party controls the evidence of its own trustworthiness. *Governance*: decisions that require capabilities controlled by external parties; the absence of internal capacity to evaluate what is governed. *Adaptation*: responses that amplify the conditions they address; adaptation that lags the rate of system change.

Proposition 9.1 (Capability Constraint)**.** *The effectiveness of any capability is constrained by the effectiveness of the capabilities on which it recursively depends. The weakest capability in the recursive chain constrains the integrity of the whole, because each capability operates on the output of those before it. Capability assessment must therefore evaluate the quality of the recursive interactions, not only the capabilities in isolation.*

10 The QCRF Assessment Scaffold

A theoretical framework that offered no means of application would be open to the charge of being purely explanatory. This section sets out an assessment scaffold. It is deliberately described as a scaffold rather than a methodology: it provides an ordinal structure an assessor can apply through reasoned judgement, while the calibration of that structure into validated, quantitative instruments is identified in Section 13 as required empirical work. The honest position is that the scaffold supports disciplined assessment now and quantification later; it does not pretend to deliver a validated metric it has not yet earned.

10.1 What Is Assessed

Assessment evaluates four dimensions: *capability* (does the institution possess the capacity to perform each governance function?); *sovereignty* (does it retain independent authority over that function?); *resilience* (can it sustain the function under disruption?); and *recursion* (does each capability strengthen or degrade the mechanisms on which the others depend?). The fourth dimension distinguishes QCRF from conventional assessment, which evaluates functions in isolation.

10.2 Capability Maturity and Sovereignty State

The two ordinal scales below are applied to each of the seven capability domains.

Maturity	Condition
1 Fragmented	Functions exist in isolation, without coordination.
2 Defined	Functions are documented and assigned.
3 Integrated	Functions interact and share evidence.
4 Sovereign	Functions operate under independent institutional authority.
5 Recursive	Functions adapt through controlled recursive feedback.

Table 2: Capability maturity scale.

State	Condition
0 None	The function is entirely externalised.
1 Partial	Authority is retained over some sub-functions.
2 Conditional	Sovereignty holds in normal conditions but degrades under stress.
3 Full	Independent authority is retained under all assessed conditions.

Table 3: Sovereignty state scale.

The two scales are orthogonal and should not be conflated, despite the coincidence of vocabulary. Maturity (Table 2) measures how well-developed and integrated a governance function is; its fourth level, "Sovereign," describes a function that operates under independent authority. Sovereignty state (Table 3) measures, separately, the degree of that independent authority on a four-point scale. A function can be highly mature in its internal development yet sit at a low sovereignty state because its critical inputs are externalised, and a function can hold full sovereignty while remaining immature in execution. The two are assessed independently and read together: maturity describes the quality of the function, sovereignty describes who controls it.

10.3 Recursive Risk Indicators

The Recursive Risk Indicators measure the health of the recursive cycle rather than the state of individual controls. Each is assessed against observable conditions rather than left as an abstraction. Table 4 states, for each indicator, the observable condition an assessor judges and the direction of concern.

Indicator	**Observable condition assessed**	**Concern when**
Measurement dependency	Proportion of decision-critical metrics defined or supplied by the measured party rather than the institution	High
Validation dependency	Whether validation can be performed without reliance on unchallengeable or proprietary mechanisms	Cannot
Assurance dependency	Whether assurance evidence originates from a source independent of the assured system	Does not
Trust concentration	Number of distinct actors or mechanisms on which decision-critical trust rests	Low
Governance delegation	Proportion of governance decisions that require externally controlled capability to execute	High
Adaptation latency	Time between a material change in system state and the institution's governed response	Long

Table 4: Recursive Risk Indicators and their observable conditions. The scaffold supports reasoned ordinal judgement; calibration into validated quantitative indices is identified as future empirical work.

These indicators provide leading signals. An institution may pass every conventional control assessment while exhibiting high measurement dependency, high governance delegation, and long adaptation latency, that is, while its recursive governance cycle is approaching instability. The scaffold is designed to make that condition visible before it produces a failure that conventional assessment would only detect afterward.

11 Quantum as the Proving Ground

The theory developed above is general. This section explains why quantum technologies are the domain in which it becomes unavoidable, and why they were the route through which the framework was discovered.

Quantum technologies do not create the need for recursive governance theory. They reveal it. Their characteristics expose, simultaneously and unambiguously, the failure modes that exist

latently in any complex dependency but are usually gradual enough to be missed.

Quantum measurement is probabilistic and may alter the state being measured, which stresses the observation and measurement layers and challenges Measurement Sovereignty directly. Quantum validation is frequently model-dependent, and some outputs may require equivalent quantum systems to check, which stresses the validation layer. Quantum assurance is often vendor-dependent, because the specialised hardware, expertise, and supply chains are held by a small number of actors, which stresses the assurance layer and drives concentration. Quantum governance becomes capability-dependent, because the institution cannot evaluate what it cannot reproduce. And quantum adaptation operates under compressed timelines, which stresses the adaptation layer.

The significance is not that any one of these is unique to quantum. Other domains have governed stochastic, complex, partially observable systems for decades, and have done so with sophisticated probabilistic methods. Nuclear safety, quantitative finance, and aerospace are the obvious examples, and the framework does not claim they lack rigorous risk governance. The distinctive feature of quantum is that it stresses all the layers at once, through a single category of capability, and externalises sovereignty at each of them simultaneously. In the established high-reliability domains, the governing institutions retained Measurement Sovereignty: the regulator could, in principle, independently measure and validate. The quantum case is distinctive because that independent capacity is, for many adopting institutions, absent from the outset.

Artificial intelligence is the domain in which the same dynamics are already operating at scale. Model outputs change without system modification; validation is model-dependent; assurance is concentrated among a small number of providers; and adaptation timelines are compressing. The structural pattern is identical to the quantum case. The difference is only that AI dependency is already widespread, whereas quantum dependency is still forming. Quantum is therefore the proving ground in two senses: it demonstrates the theory most clearly, and it offers the opportunity to apply the theory before the dependency becomes irreversible.

11.1 A Worked Illustration

A short, concrete sketch shows what the linkage looks like in practice. Consider a regulated financial institution that adopts a vendor-supplied model, quantum-inspired or AI-based, to optimise its capital allocation and price portfolio risk. The adoption is rational: the model outperforms the institution's internal methods, and competitive pressure makes its use close to mandatory.

Trace the cycle. At the *measurement* layer, the institution can no longer independently observe how the model represents the state of its portfolio; the relevant metrics are defined by the vendor's model, not by the institution. At the *validation* layer, the institution cannot independently confirm that the model's outputs are correct, because the model is proprietary and, in the quantum-inspired case, may not be reproducible on the institution's own systems. At the *assurance* layer, confidence in the model rests on the vendor's own testing and on third-party assurance that itself relies on the vendor's disclosures; the assurance is, in the sense of Section 5, circular. At the *trust* layer, reliance is now placed on an actor whose trustworthiness the institution cannot independently evidence. At the *governance* layer, the board continues to approve risk appetite and sign regulatory attestations, but the substantive determination of risk has migrated to the vendor's model. At the *adaptation* layer, when market conditions shift, the institution's capacity to adjust depends on the vendor's willingness and speed in updating the model.

Every governance activity continues. Reports are produced, approvals are given, assurances are filed. Yet by the criterion of Definition 2.4, the institution's governance has been transformed: it retains the form of control while the authority over the mechanisms that produce its risk evidence

has passed outside its boundary. If the vendor's market position is strong, as it tends to be for exactly the reasons set out in the discussion of concentration, the institution cannot readily recover that authority by switching providers. This is the dependency–sovereignty linkage in a single, recognisable case, and it is consistent with the emphasis recent financial-sector analysis has placed on vendor concentration and on the difficulty of independently verifying advanced models [13]. Figure 7 sets out the case layer by layer, showing which side of the sovereignty boundary holds authority at each.

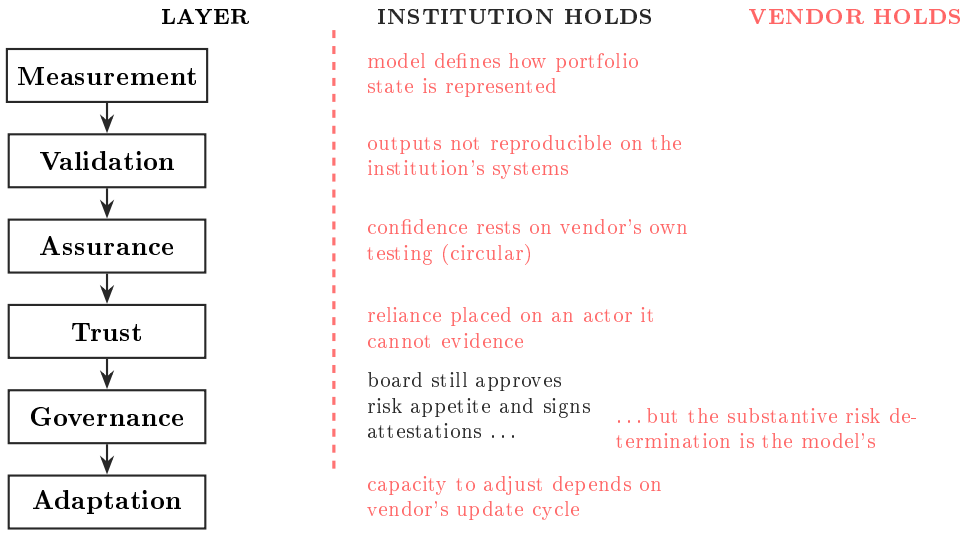

Figure 7: The worked illustration, layer by layer. A regulated institution adopts a vendor-supplied risk model. At every layer except the formal act of governance, authority sits on the vendor's side of the sovereignty boundary. The governance layer is the revealing one: the board retains the form of control, approving and attesting, while the substance of the risk determination has migrated to the model. This is Governance Transformation (Figure 3) reached through a wholly rational sequence of individually defensible decisions.

12 Theoretical Implications

Principle 12.1 (Governance Extension)**.** *The framework does not replace existing governance theory. It extends it by introducing dependency, sovereignty, and recursion as first-class analytical constructs and by showing that the recursive authority earlier theory presupposed is, under emerging-technology dependency, no longer given.*

The framework challenges six assumptions embedded in conventional governance.

Linear risk. Risk is not transferred; it is transformed. A risk transferred to a third party does not vanish. It becomes a dependency, a trust requirement, a governance obligation, and a concentration exposure. The recursive view reveals transformations that linear models record as reductions.

Static trust. Trust is reconstructed, not inherited. Trust held without reconstruction becomes assumption, and untested assumption becomes vulnerability.

Final assurance. Assurance requires assurance. The mechanisms of assurance rest on measurement, validation, and evidence subject to the same recursion as the systems they assure.

Governance stability. Governance is recursive and adapts to its own outputs. Stability is not a property of the framework but an emergent property of controlled recursion.

Compliance sufficiency. Compliance demonstrates conformity at a point. Sovereignty supports resilience under change. An institution may be fully compliant and entirely unable to detect, interpret, or respond to a condition its standard did not anticipate.

Technology-centric risk. A technology failure becomes a governance failure through its effect on measurement, validation, assurance, trust, and adaptation. The failure of a sensing system is a Measurement Sovereignty failure if the institution cannot detect it, a validation failure if it cannot determine the consequences, and a governance failure if it cannot adapt.

Thesis 12.1 (The Recursive Governance Thesis)**.** *Risk emerges through recursive interactions between observation, measurement, validation, assurance, trust, governance, and adaptation. It cannot be fully understood through analysis of technological, operational, or organisational conditions in isolation. Governance analysis must therefore prioritise recursive mechanisms over static artefacts.*

Thesis 12.2 (The Sovereignty Thesis)**.** *Governance resilience is constrained by the degree of sovereignty an institution maintains over the mechanisms through which evidence is interpreted, confidence is established, trust is formed, decisions are made, and adaptation occurs. Loss of sovereignty introduces governance asymmetry; asymmetry creates dependency; and dependency constrains the capacity to challenge assumptions, validate decisions, and adapt.*

13 Boundaries and Limitations

The framework is theoretical and requires empirical validation. Its purpose is explanatory: to provide a model for understanding how governance assumptions are stressed by emerging-technology dependency. It is honest about what it is not. It is not, in its present form, a turnkey audit methodology, and it does not provide validated quantitative metrics. The assessment scaffold of Section 10 supports disciplined ordinal judgement; it does not yet deliver calibrated indices, and claiming otherwise would reproduce precisely the formal-without-substance failure the framework warns against.

The framework does not claim that all governance is recursive. It claims that recursion becomes increasingly significant as complexity and dependency increase. In systems with stable evidence, limited interdependence, and slow change, conventional linear governance may remain adequate. The framework's value rises with interconnection, adaptivity, dependency, and the difficulty of independent observation.

Several claims are falsifiable, and stating them is part of the contribution. Recursive governance failure should produce outcomes that linear models cannot predict or explain. Measurement Sovereignty loss would be expected to precede, or at least correlate strongly with, systemic governance failure rather than follow it, recognising that in complex socio-technical settings strict causal ordering is difficult to establish. Governance Sovereignty, if it is genuinely measurable, should degrade observably in advance of governance failures. Concentration should amplify recursive fragility in a way distinguishable from its ordinary market effects. Each of these is testable, and each could falsify the framework.

The framework's central claim admits a single, sharp, predictive test, stated here so that the academic community has a clear target. The dependency–sovereignty linkage predicts that sovereignty-reducing dependency, where it is not offset by retained sovereignty at some layer, leads to Governance Transformation. This yields a falsifiable hypothesis:

> *Consider an institution whose primary operational capability exhibits high opacity, non-reproducibility, and non-substitutability, and which has not established indepen-*

> *dent measurement telemetry over that capability. The theory predicts that such an institution cannot maintain substantive governance over its strategic adaptations across a multi-year horizon. The theory is falsified if such an institution is observed to retain substantive governance over a five-year horizon without establishing any independent measurement capacity over the capability.*

To prevent the test from becoming unfalsifiable in practice, "substantive governance" must be anchored to observable, binary criteria rather than left to interpretive judgement. The following conditions are proposed as the minimum evidentiary standard. The institution can demonstrate that, at least once within any twelve-month period during the test horizon, it (a) ran an independent validation of the capability's outputs using instrumentation it controls, producing results it can present to a third party; (b) made a strategic adaptation to the capability's use, its parameters, its scope, or its continuity, that was initiated by the institution's own analysis rather than by the provider's recommendation or update cycle; and (c) documented the decision trail from its own evidence to the adaptation in a form auditable after the fact. If all three can be demonstrated, the institution retains substantive governance by the standard of this test. If none can, it does not. The criteria are deliberately binary and externally verifiable so that neither the institution nor a reviewer can relitigate the outcome by redefining "substantive" after the fact.

The hypothesis is constructed to be defeatable. It names the conditions (the three properties plus absent telemetry), the outcome (loss of substantive governance over strategic adaptation), the observable that would distinguish substantive from formal governance (the provenance and independent validation of adaptations), and the time horizon over which the prediction holds. A single well-documented counterexample, an institution meeting the antecedent conditions yet retaining genuine adaptive sovereignty over the period, would require the theory to be revised or abandoned. Establishing the prevalence of the antecedent conditions, and tracking the predicted outcome longitudinally, is the empirical programme the framework calls for.

The empirical base is currently narrow. The framework was developed through the lens of quantum-technology adoption, and although the theoretical structure is intended to be general, its empirical grounding is primarily in quantum and AI, with the application to autonomous-systems governance, synthetic biology, and other domains resting at present on conceptual argument rather than demonstration. The work is therefore properly described as a proposed general theory whose generality requires empirical testing across domains, not as a generality already established. The required work includes: formal and computational modelling of recursive governance dynamics; operationalisation and validation of the sovereignty constructs as measurable properties; a full taxonomy of sovereignty, addressing whether each kind of sovereignty can be retained, shared, delegated, federated, leased, or recovered, and under what conditions; longitudinal study of dependency formation, including the empirical prevalence of sovereignty-reducing dependency that this paper can only infer; comparative analysis across capability domains; retrospective application to documented governance failures to test explanatory power against conventional accounts; and the development of cost-proportionality guidance for independent measurement, addressing how the depth and frequency of independent verification should vary by sector, capability criticality, and governance exposure, so that the framework's requirements are calibrated to the risk rather than applied uniformly.

Proposition 13.1 (Research Proposition)**.** *The primary value of the framework lies in its capacity to expose, challenge, and re-examine governance assumptions that remain implicit within conventional models of risk, assurance, trust, and decision-making. It is offered as a contribution to governance theory and a scaffold for disciplined assessment, not as a replacement for established governance practice.*

14 Toward a General Theory of Recursive Governance

Governance becomes increasingly recursive as systems become more interconnected, more adaptive, and more dependent upon capabilities that modify the conditions under which governance operates. The constructs developed here, Recursive Trust, Recursive Assurance, Measurement Sovereignty, Governance Sovereignty, and the recursive architecture that integrates them, describe a governance model in which each function modifies the conditions for the next, and in which the institution's authority over those functions is the variable that determines resilience.

The decisive shift is from governance functions to governance relationships. Governance rarely fails because a function is absent. It fails when the relationships between functions degrade: when measurement no longer reliably informs validation, when validation no longer reliably informs assurance, and so on through the cycle. An institution may possess every function and still fail, because the recursive relationships among them have broken. Governance effectiveness is a property of relationships, not of components. This is the heart of the Recursive Governance Thesis.

The second shift is from authority assumed to authority maintained. The cybernetics tradition gave us the recursive cycle and assumed the institution controlled it. The contribution of this work is to show that the control is not given. Dependency formation erodes it, layer by layer, and sovereignty is the name for what is being lost. This is the heart of the Sovereignty Thesis, and it is what allows the framework to extend the established tradition rather than merely restate it.

Quantum technologies are the proving ground because they reveal the erosion earliest and most completely. Artificial intelligence is the domain where the erosion is already underway. Autonomous systems, synthetic biology, and whatever follows will exhibit the same structure, because the structure is a property of dependency, not of any particular technology.

The challenge for future research is not whether governance is recursive; the evidence for recursive dynamics is visible across technology-adoption cycles, regulatory responses, and assurance failures. The challenge is whether governance theory, and the institutions that rely on it, can develop the sovereignty constructs quickly enough to retain authority over their own mechanisms before dependency removes the option.

The asymmetry between loss and recovery sharpens the point. Sovereignty is lost gradually, through individually rational decisions, while governance activity continues undisturbed; it is recovered only through shock, external intervention, or a faculty deliberately kept in reserve (Section 7.4). Because recovery from total dependency cannot be generated from within the corrupted cycle, the decisive moment is not the recovery but the retention: the point at which an institution chooses, before dependency is complete, to keep independent measurement over something that matters. That choice is cheap while sovereignty remains and irrecoverable once it is gone.

The paper began with a single claim, and it ends by returning to it. The significance of a capability is not what it can achieve, but the dependencies it creates, the trust it requires, the governance it obliges, and the concentration it invites. An institution that adopts a capability without measuring it on those terms has not made a technology decision. It has made a sovereignty decision, and it has made it without knowing.

References

[1] W. R. Ashby, *An Introduction to Cybernetics*. London: Chapman & Hall, 1956.

[2] S. Beer, *Brain of the Firm: The Managerial Cybernetics of Organization*. London: Allen Lane, 1972.

[3] N. Luhmann, *Social Systems*. Stanford: Stanford University Press, 1995. Translated by J. Bednarz Jr. and D. Baecker.

[4] C. Perrow, *Normal Accidents: Living with High-Risk Technologies*. New York: Basic Books, 1984.

[5] E. Hollnagel, D. D. Woods, and N. Leveson, eds., *Resilience Engineering: Concepts and Precepts*. Aldershot: Ashgate, 2006.

[6] M. Power, *The Audit Society: Rituals of Verification*. Oxford: Oxford University Press, 1997.

[7] D. H. Meadows, *Thinking in Systems: A Primer*. White River Junction: Chelsea Green, 2008.

[8] H. R. Maturana and F. J. Varela, *Autopoiesis and Cognition: The Realization of the Living*. Dordrecht: D. Reidel, 1980.

[9] N. N. Taleb, *Antifragile: Things That Gain from Disorder*. New York: Random House, 2012.

[10] E. Ostrom, *Governing the Commons: The Evolution of Institutions for Collective Action*. Cambridge: Cambridge University Press, 1990.

[11] National Institute of Standards and Technology, "Announcing Approval of Three Federal Information Processing Standards for Post-Quantum Cryptography," Aug. 2024.

[12] National Cyber Security Centre, "Timelines for Migration to Post-Quantum Cryptography," Mar. 2025.

[13] G7 Central Bank Quantum Technologies Working Group, "Preparing for Quantum Technologies: Key Considerations for Financial Sector Participants," May 2026.

[14] World Economic Forum and Deloitte, "Transitioning to a Quantum-Secure Economy," White Paper, Sept. 2022.