Decision Without Deciders: AI Governance and the Infrastructuralization of Accountability

Abstract

AI systems are transforming how consequential decisions are produced and how accountability is organized. Rather than functioning as tools operated by identifiable agents, algorithmic systems increasingly operate as decision environments in which outcomes emerge from distributed socio-technical configurations. This paper theorizes this transformation as responsibility absorption: the structural migration of accountability from traceable relations between agents and outcomes into infrastructural architectures, producing the condition of decision without deciders. In such systems, decisions occur without identifiable decision-makers in the sense required for conventional accountability, not because responsibility disappears but because it is reconstituted as a systemic property of configuration rather than an attribute of agents.

The paper develops a structural framework for understanding this transformation. It contrasts chain-based accountability models, which organize responsibility through discrete decisional acts and identifiable agents, with field-based decision structures, in which outcomes emerge from distributed configurations. It then theorizes responsibility absorption as the displacement of agency, authority, and normative responsibility into infrastructural systems. Accountability is subsequently reconceptualized as an infrastructural property—a characteristic of system architecture rather than a regulatory overlay imposed through oversight mechanisms.

Building on this reconceptualization, the paper articulates three principles for embedding accountability within system architectures: binding authority to configuration, preserving foreclosed alternatives, and architectural contestability. Together, these principles define a model of structural placement in which responsibility is positioned within socio-technical systems rather than retrospectively attributed to individual actors.

The paper concludes that AI governance cannot be adequately addressed through technical fixes, transparency mandates, or regulatory extensions of existing accountability frameworks. Instead, it requires the construction of accountability infrastructures—institutional architectures capable of rendering distributed configurations answerable in societies increasingly organized through infrastructural decision systems. AI governance is thus reframed not as the regulation of a technology but as the architectural design of decision environments in conditions of decision without deciders.

---

1. Introduction

A physician reviews a diagnostic recommendation generated by a machine learning system. The recommendation derives from patterns identified across millions of prior cases, weighted through processes the physician cannot reconstruct and shaped by design choices made by engineers the physician will never meet. When asked who decided this patient should receive this diagnosis, no coherent answer exists. The physician invoked the system but did not author its logic. The engineers constructed the system but did not intend this specific output. The institution deployed the system but does not control its internal operations. The training data shaped the outcome but cannot be interrogated as an agent. A decision has occurred—one with profound consequences—yet the social architecture of accountability finds no stable ground.

This describes an increasingly general condition: the displacement of decision-making from discrete acts performed by identifiable agents into distributed processes embedded in infrastructural systems. AI systems do not simply obscure or distribute responsibility; they absorb it into infrastructural configurations operating according to organizational principles fundamentally different from the chain-based models of causation and attribution grounding conventional accountability structures. Dominant approaches to AI governance—ethics frameworks, regulatory proposals, transparency mandates, explainability requirements—presume the problem is one of constraint, visibility, or control applied to systems that would otherwise operate without accountability. These approaches share an assumption: accountability remains intact as a social relation requiring only extension or adaptation to incorporate algorithmic components.

Such approaches misconceive the structural transformation at stake. Responsibility does not disappear but ceases to function as a recoverable property of identifiable actors. It becomes an emergent characteristic of systemic architecture—a function of how decision-making capacity is distributed across socio-technical assemblages rather than vested in discrete agents.

We theorize this transformation through three linked concepts. *Responsibility absorption* designates the process through which accountability migrates from traceable chains of delegation into distributed infrastructural fields. *Infrastructural accountability* reframes governance as the architectural design of decision environments rather than the regulation of decision-makers. *Structural placement* identifies governance as the positioning of responsibility within socio-technical systems rather than its post hoc attribution. These concepts require reconceiving AI governance as institutional architecture—the deliberate construction of decision structures embedding accountability as a constitutional rather than regulatory property.

Situating the Argument

This paper contributes to three overlapping intellectual conversations. First, it extends work in accountability studies examining how responsibility is organized in complex institutional settings, shifting focus from individual agents to distributed organizational processes. Second, it builds on infrastructure studies within science and technology studies (STS), which has theorized how technical systems become embedded in social practices and reshape the conditions under which social life unfolds. Third, it engages structural sociology's concern with how social arrangements enable or constrain action, treating AI systems as governance architectures rather than neutral tools. Where existing work on algorithmic accountability has focused primarily on transparency, fairness, and bias, this paper reframes the problem as structural transformation in how decisions are made and responsibility is placed.

The Argument

Section 2 analyzes the structural transformation of decision-making, contrasting chain-based models presuming discrete decisional acts with field-based models treating decisions as emergent properties of systemic configurations. Section 3 develops the concept of responsibility absorption, examining how accountability becomes structurally displaced from identifiable agents into distributed infrastructures, producing *decision without deciders*. Section 4 reframes accountability as an infrastructural property rather than a moral or legal attribute, establishing the theoretical basis for architectural governance. Section 5 articulates three structural principles for infrastructural accountability: binding authority to configuration, preserving foreclosed alternatives, and establishing architectural contestability. The conclusion considers implications for governance theories adequate to socio-technical systems characterized by decision without deciders.

This paper offers neither ethical prescriptions, policy recommendations, nor technical specifications. The medical diagnostic scenario serves as a theoretical site for examining structural dynamics in a domain where responsibility has been extensively institutionalized. The theory remains domain-general, addressing transformations cutting across contexts in which algorithmic systems mediate consequential social outcomes.

---

2. Structural Transformation of Decision-Making

This section establishes the conceptual foundation for understanding how AI systems restructure decision-making. It contrasts chain-based structures organizing conventional accountability with field-based structures characterizing algorithmic systems.

2.1 Chain-Based Accountability Structures

Modern institutional accountability presumes a *chain-based* model of decision-making. Decisions exist as discrete acts performed by identifiable agents possessing both epistemic access to decision-relevant information and institutional authority to determine outcomes. Responsibility flows through traceable chains of delegation: a subordinate acts under a superior's authority; an agent operates within a principal's mandate. When outcomes require justification or redress, accountability structures trace backward through these chains to locate responsible agents.

This model depends on four structural features. *Temporal discreteness*: decisions occur as distinguishable events separable from their context, and the moment of decision can be identified. *Agential localization*: decisional authority resides in specifiable actors whose identity and institutional position can be determined. *Epistemic traceability*: the reasoning producing decisions can be reconstructed and evaluated. *Normative accountability*: agents can be held responsible because they possessed the capacity and authority to decide otherwise.

These features enable *retrospective accountability*—the ability to trace outcomes backward through decisional chains to establish responsibility. When a medical error occurs, investigators identify which clinician made which determination. When a sentence appears unjust, scrutiny focuses on the judge's reasoning. Chain-based structures make accountability operative by ensuring outcomes can be attached to decisions, decisions to decision-makers, and decision-makers to normative frameworks rendering their actions evaluable.

2.2 Field-Based Decision Structures

AI systems introduce *field-based* decision structures—configurations in which outcomes emerge from the interaction of distributed components rather than from discrete acts by identifiable agents. Decisions cannot be localized to specific moments or actors. Instead, they are systemic properties produced through the mutual constitution of heterogeneous elements: algorithmic models, training datasets, architectural designs, institutional deployment contexts, user interactions, feedback loops, and regulatory environments.

The diagnostic system's recommendation emerges from the selection of training data, which determines what patterns the system can recognize; the choice of model architecture, which constrains what relationships can be learned; the specification of optimization objectives, which defines what counts as performance; the institutional context of deployment, which shapes how outputs are interpreted; the interface design, which structures user engagement; and feedback mechanisms, which enable or prevent learning. Each component contributes to system behavior, yet none individually determines specific outputs. The recommendation represents not a decision made by a decider but an emergent property of systemic configuration.

Field-based structures exhibit characteristics distinguishing them from chain-based models. *Configurational determination*: outcomes derive from how components are arranged relative to one another rather than from individual component properties. The same algorithmic model produces different outputs when trained on different data or deployed in different institutional contexts. Decision-making capacity resides in *configuration*—the systematic interrelation of elements.

Distributed temporality: the "moment of decision" cannot be temporally localized. When did the system decide to generate this recommendation—when training data was selected, when the model was trained, or when the institution deployed it? Each moment contributes; none alone constitutes "the decision."

Emergent determination: outcomes cannot be predicted or controlled through modifying individual components because they emerge from component interactions. The relationship between component modification and systemic outcomes is non-linear and often opaque.

Operational opacity: the processes generating outcomes become epistemically inaccessible not merely due to technical complexity but due to structural features of how field-based systems operate. The recommendation emerges from pattern identification across millions of cases, weighted through mathematical operations lacking semantic content. Opacity is not a correctable deficiency but a structural feature.

2.3 From Decisions as Acts to Decisions as Environments

The shift from chain-based to field-based structures represents the *environmentalization of decision-making*. In chain-based structures, decisions constitute discrete acts occurring *within* environments—contexts that constrain but remain analytically separable from decisional acts. A judge decides within an institutional environment of legal precedent and professional norms, but the decision itself remains distinguishable. In field-based structures, decisions *are* environments—systemic configurations shaping outcomes without being reducible to discrete acts by identifiable agents.

This transformation generates the *accountability gap*—a structural disjunction between how responsibility is produced (through distributed field-based processes) and how it is institutionally recognized (through chain-based attribution structures). Existing accountability frameworks seek to bridge this gap by attempting to localize responsibility within field-based structures: holding deploying institutions liable, requiring human oversight, and mandating explainability. These approaches impose chain-based accountability structures onto field-based decision processes.

Such strategies misconceive the structural transformation at stake. The accountability gap emerges not from insufficient regulation or inadequate transparency but from the displacement of decision-making from chains to fields—from discrete acts to distributed configurations. Addressing this transformation requires not adapting chain-based structures to incorporate algorithmic components but developing accountability frameworks adequate to field-based decision environments. This requires theorizing how to make field-based structures themselves accountable as infrastructural configurations.

---

3. Responsibility Absorption

This section develops the core concept of responsibility absorption and examines its implications, showing how accountability becomes structurally displaced from agents into systems, producing decision without deciders.

3.1 The Concept of Responsibility Absorption

Responsibility absorption is the process through which accountability ceases to function as a recoverable attribute of identifiable agents and becomes an emergent characteristic of systemic configuration. Absorption does not imply elimination. Rather, responsibility undergoes an ontological shift: it migrates from a social relation between actors and outcomes into a structural property of socio-technical assemblages. Responsibility is absorbed into architecture insofar as accountability can no longer be extracted and reassigned to specific agents through conventional institutional mechanisms.

Responsibility absorption operates through four mechanisms. *Distributional diffusion*: decisional influence disperses across actors and processes such that no single site possesses sufficient control to be designated “the decider.” *Temporal extension*: responsibility stretches across time horizons that exceed the temporal scope of conventional accountability practices. Configuration choices generate consequences long after designers have moved on and in contexts never anticipated. *Normative displacement*: the criteria that generate outcomes become embedded in technical specifications that resist evaluation through conventional normative vocabularies. Value judgments about what counts as a good outcome are operationalized in mathematical terms that present as value-neutral. *Operational autonomization*: systems exhibit behaviors not directly authored by any agent. Machine-learning models extract patterns and generate outputs that exceed designers' intentions and deployers' expectations.

3.2 Decision Without Deciders

Responsibility absorption produces *decision without deciders*—a condition in which consequential outcomes occur through processes that lack identifiable authors in the sense required for conventional accountability. This does not mean AI systems possess agency. It means the social architectures that have rendered decisions accountable—structures that locate authority, knowledge, and normative responsibility in identifiable agents—no longer gain traction on how outcomes are actually produced.

The diagnostic system generates recommendations that shape clinical practice, yet no one straightforwardly "decided" what the system would recommend. Engineers selected architectures but did not determine particular outputs. Institutions chose deployment but did not author the system's internal logic. Clinicians chose to consult the system but did not construct its knowledge base. A recommendation emerges—one that may determine whether a patient receives life-saving treatment—yet accountability practices find no stable ground.

Decision without deciders differs from familiar forms of distributed decision-making in complex organizations. When a board approves a merger, responsibility may be collective, but it remains locatable: the deciding body can be identified, its deliberation reconstructed, and its reasoning evaluated. By contrast, decision without deciders describes outcomes that emerge from configurations that resist reduction to the decisions of any identifiable agent or collective body. No entity possesses the combination of epistemic access, institutional authority, and operational control that would position it as "the decider."

This condition poses a direct challenge to institutional legitimacy. Modern governance practices legitimate consequential outcomes by rendering them accountable to those affected. Chain-based accountability accomplishes this by producing determinate sites where authority is exercised and responsibility attaches. Decision without deciders disrupts that architecture. When outcomes are produced by infrastructural configurations rather than identifiable agents, the mechanisms through which institutions establish legitimacy lose their structural footing.

3.3 The Structural Displacement of Agency

Responsibility absorption does not eliminate human agency; it relocates it within socio-technical systems. Agency becomes *structurally displaced*: it shifts away from the point where outcomes are generated and toward sites where architectures are configured. Meaningful discretion migrates from operators to those who design, train, deploy, and maintain systems. Yet this displacement does not simply relocate responsibility in a form that chain-based accountability can accommodate. Configuration itself is typically distributed, performed by multiple actors under partial knowledge and constraint.

Consider training data selection—a configuration choice with profound implications for system behavior. Curators make explicit judgments, but those judgments are shaped by data availability, institutional priorities, resource limitations, and technical requirements. The "decision" about training data is therefore not a single act but an interaction among curatorial discretion, organizational constraint, technical possibility, and historical contingency. No single agent exercises comprehensive authority, yet the configuration decisively shapes system behavior.

The structural displacement of agency means responsibility for outcomes cannot be cleanly attached to operators (who lack control over system behavior), designers (who lack knowledge of deployment contexts), or deployers (who lack control over internal operations). Responsibility becomes structurally indeterminate—not because it is hidden but because it is constituted as a systemic property rather than an attribute of agents. This indeterminacy is not resolved by improving traceability; it is a structural condition that requires accountability frameworks adequate to infrastructural decision environments.

3.4 Implications for Governance

Responsibility absorption shows that conventional governance responses—those premised on identifying responsible agents and establishing oversight—address symptoms rather than structure. Transparency mandates assume that making system operations visible will enable responsibility attribution. Yet under responsibility absorption, accountability has already migrated from identifiable agents into configurations. Visibility does not restore chain-based accountability because the core problem is not secrecy but structural transformation.

Similarly, "human in the loop" governance presumes that inserting human judgment preserves chain-based responsibility. But when systems function as decision environments rather than tools, oversight does not restore agency. The human inherits a decision environment already structured by training data, optimization objectives, interface design, and institutional thresholds. Human review operates within an infrastructural field that shapes what appears as a decision and what options are available.

Governance adequate to responsibility absorption cannot proceed by searching for responsible agents within field-based systems. It requires reconceiving accountability as an infrastructural property—a characteristic of how systems are architected rather than an attribute of operators. This shifts the focus of governance from regulating decision-makers to designing decision environments, from post hoc attribution to structural placement, and from oversight to architecture.

---

4. Accountability as Infrastructure

This section reframes accountability from a regulatory requirement to an infrastructural property, showing how governance must shift from oversight to architectural design.

4.1 Beyond Regulatory Accountability

Dominant approaches conceptualize accountability as *regulatory accountability*—oversight mechanisms, compliance requirements, and redress procedures applied to existing systems.

Regulatory accountability treats accountability as external to systems, imposed through rules that constrain or evaluate system behavior. This presumes that systems, left unregulated, would operate without accountability.

This misconceives the relationship between systems and accountability. When decision-making becomes infrastructural—when outcomes emerge from systemic configurations—accountability cannot be externally imposed but must instead be constitutive of system architecture itself. Accountability thus becomes an *infrastructural property*: a characteristic of how systems are designed, configured, and operated, rather than a regulatory overlay applied post hoc.

Infrastructural accountability reconceives governance as a problem of *architectural design*—how decision environments are constructed rather than how decisions are evaluated. If outcomes emerge from systemic configurations, ensuring accountability requires the deliberate architectural design of those configurations so that responsibility becomes traceable, contestable, and operationalizable.

This reframing has three implications.

First, it shifts the temporal focus from *ex post* evaluation to *ex ante* design. Regulatory accountability emphasizes auditing outputs, investigating errors, and providing redress. Infrastructural accountability emphasizes architectural decisions during development: how training data is selected, how model architectures are specified, and how outputs are interpreted. The critical moments for governance thus become those in which system configurations are established.

Second, it shifts the locus of governance from oversight bodies to design processes. Regulatory accountability vests authority in external entities—regulators, auditors, and review boards. Infrastructural accountability locates governance within the processes through which systems are built: development practices, deployment decisions, and institutional arrangements. Governance therefore becomes distributed across the sites where configuration decisions are made.

Third, it shifts the mode of governance from constraint to construction. Regulatory accountability operates through prohibition: systems must not do certain things. Infrastructural accountability operates through architectural specification: systems must be built such that certain accountability properties are necessarily present. Rather than constraining what systems can do, infrastructural accountability shapes what systems *are*.

4.2 Responsibility as Architectural Property

If accountability is infrastructural, responsibility must be reconceived not as an attribute of agents but as an *architectural property* of socio-technical systems—a characteristic emerging from how systems are configured rather than from individual actor intentions.

An architectural property inheres in structural relations rather than component characteristics. The load-bearing capacity of a bridge emerges from how components are arranged relative to one another; the same steel beams produce different load capacities depending on structural configuration. Similarly, the accountability of a decision-making system is architectural: it emerges from how decisional components—algorithms, data, interfaces, institutional contexts, and human actors—are arranged relative to one another.

Treating responsibility as architectural means accountability is not something individual agents possess but something systemic configurations enable or preclude. A physician operating within a system that provides no access to the reasoning behind diagnostic recommendations cannot be held responsible in the manner conventional medical accountability requires, regardless of competence. The system architecture renders the physician structurally incapable of the epistemic engagement medical responsibility presumes. Conversely, a system architected to make reasoning accessible and outputs provisional enables accountability even when decision-making is algorithmically mediated.

This architectural conception reveals that responsibility is *designed into* systems rather than *attributed to* actors. Systems that distribute decision-making across opaque algorithmic processes and inaccessible training data produce architectures that structurally preclude certain accountability forms. Systems that expose reasoning and preserve alternatives produce architectures that enable accountability. The presence or absence of accountability is therefore not a matter of individual agent behavior but of how systems are constituted.

Responsibility as architectural property means governance must focus on *structural placement*—the deliberate positioning of accountability-enabling features within system architectures. AI governance must position accountability within socio-technical systems to shape how responsibility operates. This positioning occurs through design choices: where in development pipelines value judgments are made explicit, where opportunities for contestation are embedded, where alternatives are preserved, and where authority to intervene is located.

4.3 Accountability Architecture and Institutional Design

Infrastructural accountability requires reconceiving AI governance as a problem of *institutional design*. If accountability is architectural, building accountable systems means constructing institutional arrangements that embed responsibility into socio-technical architectures. This is not primarily a technical problem—it does not turn on developing better algorithms or achieving greater transparency. It is an institutional problem: how social relations, authority structures, knowledge practices, and normative frameworks are organized in the development and deployment of AI systems such that accountability becomes constitutive of system operation.

Infrastructural accountability therefore requires institutional innovations that enable responsibility to function as an architectural property: governance structures positioning diverse stakeholders in configuration decisions; knowledge practices making the normative implications of technical choices explicit; authority arrangements aligning decisional influence with accountability capacity; interface designs preserving human engagement; organizational cultures treating accountability as constitutive rather than compliance-oriented; and legal frameworks recognizing systemic configuration as a proper object of governance.

The institutional character of infrastructural accountability means governance cannot be accomplished through technical fixes alone. No amount of algorithmic transparency can substitute for institutional arrangements that embed responsibility into system architectures. Conversely, appropriate institutional arrangements can enable accountability even when systems are technically complex or opaque, by positioning actors with the knowledge and authority required to make and answer for configuration decisions. The challenge is not making algorithms accountable but making the socio-technical assemblages within which algorithms operate accountable—an institutional rather than technical problem.

---

5. Structural Placement Principles

This section articulates three principles for positioning accountability within system architectures. These principles specify how responsibility can be embedded as an infrastructural property.

5.1 Binding Authority to Configuration

The first principle is *binding authority to configuration*: ensuring that those who exercise decisional influence through system design possess institutional authority and normative responsibility commensurate with that influence. This responds to the structural displacement of agency from operation to configuration. If meaningful discretion resides in configuration decisions—choices concerning training data, model architectures, optimization objectives, and deployment contexts—then accountability requires that such discretion be exercised by actors positioned to answer for its consequences.

Binding authority to configuration means establishing institutional structures that recognize configuration decisions as consequential acts requiring deliberation, authorization, and justification. Currently, many configuration decisions are made by technical specialists operating under professional norms of engineering optimization rather than normative frameworks of social consequence. Engineers select training data based on availability and technical suitability, and specify optimization objectives based on measurable performance, without necessarily recognizing these choices as embodying normative commitments. Configuration thus occurs as technical implementation rather than as institutional decision-making.

Infrastructural accountability requires transforming configuration into a governed process—one subject to forms of deliberation, authorization, and accountability appropriate to consequential institutional decisions. Several mechanisms may operationalize this principle:

Institutional authorization requirements: consequential configuration decisions require explicit authorization from institutional bodies positioned to evaluate normative implications. Such bodies do not make technical decisions but assess whether proposed configurations satisfy accountability conditions.

Multi-stakeholder governance structures: configuration processes include representatives of affected communities, domain experts, and accountability specialists alongside technical developers, ensuring plural perspectives in architectural formation.

Normative specification requirements: configuration decisions require explicit articulation of the normative commitments embedded in technical choices. When optimization objectives are specified, the values they operationalize must be made explicit. When training data is selected, the representational scope of experience must be articulated.

Accountability roles: organizations establish positions specifically responsible for accountability dimensions of system development, ensuring the systematic integration of accountability concerns throughout configuration processes.

This principle recognizes that responsibility for outcomes cannot be separated from control over the configurations that produce them. If engineers exercise decisive influence over system behavior through configuration decisions, they must be institutionally recognized as exercising discretion requiring accountability. Conversely, if institutions seek to be accountable for system outcomes, they must exercise authority over configuration decisions. Binding authority to configuration aligns institutional structures with the actual distribution of decisional influence in AI systems.

5.2 Preserving Foreclosed Alternatives

The second principle is *preserving foreclosed alternatives*: maintaining systematic records of configuration options considered, choices made, and rationales provided. This responds to the temporal extension of responsibility in AI systems, where configuration decisions made at one moment generate consequences across extended temporal horizons and unanticipated contexts. Preserving alternatives enables both retrospective accountability and prospective adaptation.

Foreclosure of alternatives occurs throughout system development. When training data is selected, alternative datasets are rejected. When model architectures are chosen, alternative designs are foregone. When optimization objectives are specified, alternative metrics are passed over. Each choice forecloses paths that might have shaped system behavior differently. In chain-based structures, such foreclosures are institutionally visible—policy decisions are documented, judicial reasoning is recorded. In field-based structures, foreclosures often occur without systematic documentation, rendering reconstruction of alternatives impossible.

Preserving foreclosed alternatives requires institutional practices and technical systems that systematically document configuration decisions and their contexts. This preservation serves multiple accountability functions:

Retrospective intelligibility: when problematic outcomes emerge, preserved alternatives enable analysis of which configuration choices shaped those outcomes.

Justificatory accountability: preserved alternatives allow actors to justify configuration decisions by articulating what options existed and why particular paths were selected.

Adaptive capacity: preserved alternatives enable revision of system architectures when conditions change or new knowledge emerges.

Contestatory resources: preserved alternatives provide materials for contestation by affected actors, enabling evaluation of whether different configurations would better serve their interests.

Preservation of alternatives resolves a central tension in infrastructural accountability: the need for stability and revisability. Systems must remain operationally stable, yet architecturally revisable. Preserved alternatives allow architectures to remain fixed for functional purposes while remaining institutionally revisable as conditions evolve. Configuration decisions thus become durable without becoming irreversible.

5.3 Architectural Contestability

The third principle is *architectural contestability*: embedding mechanisms through which system configurations can be challenged, scrutinized, and revised by those affected by system outcomes. Configuration decisions embody normative commitments—about whose experiences matter, which outcomes count as success, and whose interests are prioritized—yet these commitments are often invisible or inaccessible. Contestability renders system architectures answerable to those they govern.

Architectural contestability differs from conventional outcome-focused accountability. Existing frameworks typically enable challenges to particular decisions—appeals of credit denials or diagnostic conclusions—while leaving the system architectures that generate those decisions beyond contestation. Outcome-level accountability allows contesting outputs but not the architectures that structure outputs.

Infrastructural accountability requires architectures themselves to be contestable because architectures constitute governance arrangements. Optimization objectives prioritize values. Training data encodes social representation. Interface designs structure relations between human and algorithmic judgment. These are not neutral technical choices but decisions about how authority and decision-making are organized. Contestability ensures that such choices remain subject to challenge.

Architectural contestability requires several enabling conditions:

Accessibility of configuration information: consequential architectural choices must be legible to those who may challenge them.

Standing to contest: those affected by system operations must possess recognized institutional standing to challenge architectures independent of individualized harm.

Venues for contestation: institutional spaces must exist where architectural challenges can be raised, evaluated, and acted upon.

Responsive revision capacity: organizations must possess both technical and institutional capacity to revise architectures in response to legitimate challenges.

Burden of justification: the burden shifts from challengers to architects, requiring justification of configurations against available alternatives.

Architectural contestability establishes governance as an ongoing social process rather than a one-time design act. System architectures become continuously answerable to those they affect. This does not guarantee optimal outcomes or prevent failure, but it institutionalizes accountability where it structurally matters: at the level of configuration rather than output.

The three principles—binding authority to configuration, preserving foreclosed alternatives, and architectural contestability—constitute the *structural placement* of accountability.

Together, they position responsibility within system architectures such that accountability functions as an infrastructural property rather than a regulatory overlay. They do not prescribe specific system designs but specify the structural conditions that any infrastructurally accountable system must satisfy.

---

6. Conclusion: Toward an Accountability Infrastructure for AI Societies

AI governance represents a structural transformation in how consequential decisions are made and how accountability is organized. AI systems effect *responsibility absorption*—the migration of accountability from traceable social relations between agents and outcomes into distributed infrastructural configurations, producing *decision without deciders*. Conventional governance responses—ethics frameworks, regulatory oversight, transparency mandates, and explainability requirements—presume that accountability remains fundamentally intact as a social relation, requiring only extension or adaptation. This paper has argued that responsibility absorption represents not a gap in existing accountability structures but their displacement by a different organizational form that requires frameworks adequate to field-based decision environments.

Implications and Open Questions

Several implications follow.

First, AI governance cannot be accomplished through technical solutions alone. The challenge is not making algorithms accountable but constructing socio-technical assemblages within which accountability functions as an operational property. Accountability becomes an institutional problem rather than a technical one.

Second, governance adequate to responsibility absorption requires reconceiving the temporal focus of accountability. Conventional frameworks emphasize *ex post* evaluation—audits, investigations, and redress. Infrastructural accountability emphasizes *ex ante* design—the configuration decisions that establish how systems will operate. The critical moments for governance are not when outcomes are evaluated but when architectures are established.

Third, infrastructural accountability distributes governance across the sites where configuration occurs rather than concentrating it in specialized oversight bodies. Every actor involved in system development—designers, engineers, data curators, deployers, institutional decision-makers—exercises influence over system behavior and must be structurally positioned as accountable for that influence.

Fourth, infrastructural accountability renders system architectures themselves proper objects of contestation and revision. Governance is not achieved by getting architectures right at the moment of design but by making them continuously answerable to scrutiny and challenge. Systems must be built not only to function but to be contestable—to enable those they affect to question their constitution and demand revision when conditions change or inadequacies emerge.

Open questions remain concerning how structural principles can be operationalized in specific institutional contexts, how competing accountability demands can be balanced, how stability and revisability can be simultaneously maintained, and how contestability mechanisms can be designed to genuinely empower rather than merely legitimate. These questions require continued theoretical development, empirical investigation, and institutional experimentation.

Toward Accountability Infrastructures

The stakes extend beyond AI governance to the organization of accountability in societies increasingly structured through infrastructural systems. As consequential decisions migrate from identifiable agents into distributed configurations—not only in AI but in algorithmic platforms, data systems, and networked infrastructures—the chain-based accountability structures organizing modern governance lose their structural foundation.

Either societies develop accountability frameworks adequate to infrastructural decision-making, or they confront power exercised without responsibility—not because actors are unethical or unregulated but because institutional structures for rendering power accountable no longer achieve purchase on how power operates.

Decision without deciders names a governance condition requiring not technical fixes or regulatory overlays but institutional reconstruction. Accountability must be rebuilt as infrastructure. Societies must construct institutional architectures capable of rendering distributed configurations answerable in ways discrete agents have been answerable under chain-based frameworks.

This paper has approached AI governance not as the regulation of a new technology but as the architectural design of decision environments. The transformation is already underway. The task ahead is not to restore past accountability forms but to build accountability infrastructures capable of sustaining responsibility in societies where decisions increasingly occur without deciders.