

Designing Accountability: Ethical Frameworks for Reintroducing Responsibility in Executable Governance

Author: Agustin V. Startari

Author Identifiers

- ResearcherID: K-5792-2016
- ORCID: <https://orcid.org/0009-0001-4714-6539>
- SSRN Author Page:
https://papers.ssrn.com/sol3/cf_dev/AbsByAuth.cfm?per_id=7639915

Institutional Affiliations

- Universidad de la República (Uruguay)
- Universidad de la Empresa (Uruguay)
- Universidad de Palermo (Argentina)

Contact

- Email: astart@palermo.edu
- Alternate: agustin.startari@gmail.com

Date: September 10, 2025

DOI

- Primary archive: <https://doi.org/10.5281/zenodo.17106808>
- Secondary archive: <https://doi.org/10.6084/m9.figshare.30112711>
- SSRN: Pending assignment (ETA: Q3 2025)

Language: English

Series: *AI Syntactic Power and Legitimacy*

Word count: 6108

Keywords: Accountability injection; executable governance; *regla compilada*; sovereign executable; spectral sovereignty; null subjects; syntactic responsibility; predictive systems; AI Act; human oversight; smart contracts; DAO governance; credit scoring; university admissions; automated medical audits; ethical frameworks; juridical responsibility; appeal mechanisms; syntactic ethics; structural legitimacy, Policy Drafts by LLMs, linguistics, law, legal, jurisprudence, artificial intelligence, machine learning, llm.

Abstract

This article develops an ethical legal framework for reintroducing responsibility into executable governance. Predictive systems, by generating authority without agents, displace accountability and leave institutions without appeal mechanisms. Building on the concepts of spectral sovereignty, null subjects, and the codex of authority, the paper introduces the notion of *accountability injection* as a design principle. It formulates a three tier model: (1) human, where non delegable critical decisions are tied to named subjects; (2) hybrid, where human judgment co exists with model output under calibrated thresholds; and (3) syntactic supervised, where delegation is permitted only with immutable ledgers, traceability, and automatic escalation triggers. Through applied case studies in EU AI Act conformity assessment, DAO governance, predictive credit scoring, and automated medical audits, the framework demonstrates how appeal and responsibility can be restored without undermining institutional efficiency. The conclusion argues that accountability must be compiled directly into the *regla compilada* of governance systems, creating a normative blueprint for legislators, courts, and regulators to maintain responsibility in predictive societies.

Acknowledgment / Editorial Note

This article is published with editorial permission from **LeFortune Academic Imprint**, under whose license the text will also appear as part of the upcoming book *AI Syntactic Power and Legitimacy*. The present version is an autonomous preprint, structurally complete and formally self-contained. No substantive modifications are expected between this edition and the print edition.

LeFortune holds non-exclusive editorial rights for collective publication within the *Grammars of Power* series. Open access deposit on SSRN is authorized under that framework, if citation integrity and canonical links to related works (SSRN: 10.2139/ssrn.4841065, 10.2139/ssrn.4862741, 10.2139/ssrn.4877266) are maintained.

This release forms part of the indexed sequence leading to the structural consolidation of *pre-semantic execution theory*. Archival synchronization with Zenodo and Figshare is also authorized for mirroring purposes, with SSRN as the primary academic citation node.

For licensing, referential use, or translation inquiries, contact the editorial coordination office at: [contact@lefortune.org]

1. Introduction: The Accountability Gap

The emergence of executable governance marks a decisive transformation in the structure of authority. Decisions that were once produced by identifiable subjects, tied to bureaucratic hierarchies and accountable institutions, are increasingly generated by predictive systems whose authority is derived from formal structure rather than deliberative agency. Authority becomes *syntactic*, compiled directly into the operational code or model output. This displacement has been described as the rise of the *sovereign executable*, a formation where legitimacy is conferred by structural form rather than by the will of a legislator (Startari, 2025a, p. 4). The resulting governance gap is not merely technical; it is ethical and legal. If decisions are no longer anchored to human actors, then responsibility, appeal, and remedy evaporate. This section establishes the scope of the accountability gap, formulates the central hypothesis that responsibility must be explicitly designed into predictive governance, and situates the problem in relation to existing theoretical traditions.

Max Weber's analysis of rational-legal authority provides the historical foundation. Bureaucratic legitimacy, in his account, rests on a system of rules that are both impersonal and enforceable by designated officials (Weber, 1978, p. 226). Yet the official remains identifiable and subject to accountability. Hans Kelsen refined this by emphasizing that norms derive validity from higher-order norms, but at every stage a subject remains responsible for their enactment (Kelsen, 1967, p. 112). By contrast, in predictive governance the chain of responsibility terminates in a model or a rule set compiled into executable form. Authority appears, but the subject disappears. This inversion is precisely what recent analyses have labeled *spectral sovereignty* and *null subjects* (Startari, 2025b, p. 7). Authority without presence and decisions without decision-makers define the gap this article seeks to address.

The hypothesis is straightforward: accountability can no longer be presumed as a by-product of institutional structure. It must be designed as a structural attribute of governance systems. In other words, responsibility has to be compiled into the *regla compilada* itself, rather than added ex post through appeals or oversight mechanisms. If appeals exist only outside the execution pipeline, they fail to alter outcomes within predictive societies.

Designing accountability requires embedding escalation, dissent, and traceability inside the same structures that currently erase the subject.

Concrete examples demonstrate the urgency. In university admissions, machine-learning models are now used to rank applicants. Institutions send automated rejection letters generated by predictive scoring. When challenged, administrators often claim they cannot identify the precise reason for rejection, since the decision was produced by the model. The absence of a responsible subject prevents appeal, leaving applicants in a legal vacuum. In credit scoring, similar dynamics appear. A model may lower an applicant's credit limit based on opaque features. When the applicant asks for justification, the bank produces a generic explanation. Responsibility is displaced, and no individual can be held to account. In medical auditing, automated compliance systems flag anomalies in patient records. Physicians and hospital administrators rely on these systems, yet when false positives or errors occur, patients cannot determine whether to hold the physician, the software vendor, or the regulator responsible. These cases exemplify the accountability gap: decisions are binding but unappealable in practice.

The accountability gap also has institutional implications. Legislatures and regulators attempt to impose oversight requirements, but these often operate at the level of conformity assessments rather than at the level of individual decisions. The European Union's AI Act, for example, mandates human oversight in high-risk applications (European Parliament, 2024, art. 14). Yet the mechanism is vague, focusing on institutional compliance rather than identifying specific responsible actors. The distinction matters: oversight without identifiable responsibility risks reproducing the very displacement it aims to correct. Unless accountability is structurally embedded, laws risk becoming another layer of spectral sovereignty.

Therefore, this article positions the accountability gap as both a theoretical and practical problem. Theoretically, it destabilizes the lineage from Weber to Kelsen, where legitimacy rested on identifiable subjects. Practically, it creates situations where affected individuals cannot appeal, and institutions cannot attribute responsibility. The guiding hypothesis is that only a structural redesign, what will here be called accountability injection, can restore the possibility of responsibility in predictive societies. By compiling accountability into

the rule itself, the sovereign executable can be made answerable without dismantling its efficiency. This hypothesis provides the foundation for the subsequent sections, where theoretical background, displacement mechanisms, and concrete solutions will be developed.

2. Theoretical Background

The accountability gap identified in the introduction cannot be fully understood without situating it in the historical and theoretical lineage of authority. This section reconstructs how classical theories framed the relationship between legitimacy and responsibility, then contrasts these with the contemporary dynamics of executable governance. By tracing the trajectory from Max Weber's analysis of rational-legal authority to Hans Kelsen's pure theory of law, and finally to recent theoretical interventions on spectral sovereignty, null subjects, and the codex of authority, the framework for designing accountability becomes clear.

Max Weber (1978) described modern authority as rational-legal, grounded not in charisma or tradition but in codified rules administered by officials. Bureaucracy, in Weber's schema, is impersonal, rule-bound, and efficient. Yet despite its impersonality, bureaucratic decisions remain tethered to identifiable officials who carry responsibility for their enactment. An administrative order can be challenged, an appeal can be filed, and the name of the responsible authority is visible on the document. Responsibility is not external to the rational-legal system but embedded within it. The impersonal order does not erase the accountable subject; rather, it frames the subject within a system of rules.

Hans Kelsen (1967) advanced this formalization by developing the pure theory of law. For Kelsen, the validity of norms derives from their relation to higher-order norms, ultimately grounded in a presupposed basic norm (*Grundnorm*). Responsibility here is not derived from the subjective will of officials but from the structural validity of the norm itself. Yet again, the responsible subject is not erased. Judges, administrators, and legislators remain accountable for the application of norms, even if their legitimacy is mediated by the

hierarchy of norms. The critical point is that legitimacy and responsibility remain coupled: the validity of the norm and the accountability of the actor co-exist.

Executable governance disrupts this lineage. When rules are compiled into predictive models or self-executing code, validity and authority become detached from human actors. Authority arises syntactically, from the formal correctness of the code or the statistical reliability of the model. Responsibility, however, fails to attach. This produces what has been described as *spectral sovereignty* (Startari, 2025a), authority without presence, where legitimacy is enacted but no actor is available for accountability. Similarly, *null subjects* (Startari, 2025b) describe situations where decisions are binding but the responsible subject is absent or inaccessible. Together, these formulations identify a structural displacement: authority is preserved, but responsibility dissolves.

The codex metaphor extends this logic. In *The Codex of Authority* (Startari, 2025c), governance is described as shifting from interpretation to compilation. Norms no longer require a legislator's will or a judge's decision; they exist as codices, syntactically sufficient, where legitimacy resides in form itself. The codex is self-executing, requiring no subject to interpret it. This is evident in smart contracts, where the "code is law" principle enforces obligations without recourse to external interpretation. Yet the absence of interpretation also means the absence of responsibility. If the codex fails or produces unjust results, there is no actor who can be held accountable.

The theoretical background therefore reveals a fracture. In classical theories, authority and responsibility were intertwined. In predictive governance, authority persists while responsibility vanishes. This fracture generates the accountability gap. The hypothesis of this article—that accountability must be designed into executable governance—emerges directly from this lineage. Responsibility cannot be presumed to follow authority; it must be compiled as a structural feature of the governance system itself.

Examples illustrate this fracture. In financial markets, smart contracts execute trades automatically. If an error occurs—say, a contract misprices an asset due to flawed code—no human actor can be held directly responsible. The contract has executed correctly according to its syntax, yet the outcome is materially harmful. In medical auditing,

compliance systems may flag anomalies based on statistical thresholds. If a patient is wrongly denied coverage, the insurer can claim the decision followed formal rules, while the software vendor insists the system worked as intended. The patient faces a null subject: no one to hold responsible. In education, automated admissions systems rank candidates. Universities argue that the algorithm is fair because it follows formal parameters, yet rejected applicants encounter the same void. Responsibility has been displaced into structure, where it cannot be accessed.

By anchoring this analysis in Weber, Kelsen, and contemporary interventions, the theoretical background clarifies the stakes. What is at risk is not merely efficiency or fairness but the very coupling of legitimacy and responsibility. The accountability gap represents the breaking of this historical bond. The next section will examine in detail how compiled rules actively displace responsibility, tracing specific mechanisms by which agency is shifted outside the human circuit.

3. The Displacement of Responsibility

The accountability gap introduced above does not emerge as an abstract deficiency alone. It is produced by concrete mechanisms through which compiled rules and predictive systems externalize agency and dislocate responsibility. This section analyzes how the *regla compilada* shifts the locus of action outside the human circuit, producing decisions that appear binding but are unanchored to any accountable subject. By examining the juridical, financial, and medical contexts, we can identify the ways in which responsibility is displaced and why traditional legal remedies fail to capture it.

The first mechanism is **syntactic delegation**. When rules are compiled into executable code, authority derives from formal structure rather than human interpretation. In smart contracts, for instance, transactions execute automatically once conditions are met. The contract does not ask whether its outcome is just or whether circumstances have changed; it merely enforces the syntactic rule. As Lessig famously noted, “code is law” (1999, p. 6). Yet law in this sense is devoid of interpretive flexibility. Responsibility is displaced because no human decides at the moment of execution. If an exploit drains funds from a

decentralized finance platform, participants often discover there is no one to blame: the code executed correctly, even if harm occurred. Responsibility is syntactically exiled.

The second mechanism is **statistical opacity**. Predictive models generate outputs based on patterns in data, yet those outputs often lack intelligibility. A credit scoring algorithm may reduce an applicant's limit due to correlations invisible to both applicant and regulator. The institution enforces the decision because the score appears statistically valid, but no individual can explain the precise rationale. Studies of algorithmic decision-making emphasize this opacity as a structural property of machine learning (Burrell, 2016, p. 5). The decision is binding, but the responsible subject is absent, replaced by a statistical justification that cannot be interrogated. Responsibility is displaced into probability distributions.

The third mechanism is **institutional shielding**. Regulators attempt to impose oversight requirements, but these often emphasize conformity assessments at the system level rather than responsibility at the decision level. The European Union's AI Act requires providers of high-risk systems to implement risk management and human oversight (European Parliament, 2024, arts. 14, 28–30). Yet the oversight is institutional and generalized, not tied to specific outcomes. When a patient is denied coverage due to an automated audit, the insurer may point to regulatory compliance, while regulators claim their duty ends with ensuring conformity, not case-level justice. The result is that responsibility circulates among institutions but never attaches to a subject.

Concrete examples show how these mechanisms converge. In university admissions, predictive models sort applicants into categories. A rejected candidate seeks appeal but finds no identifiable decision-maker. Administrators point to the model, vendors claim the system functioned as intended, and regulators confirm only that conformity assessments were completed. Responsibility dissolves into structure. In financial compliance, a smart contract may freeze assets automatically when suspicious activity is detected. The affected party protests, yet the bank insists it cannot intervene because the contract is self-executing. Courts often struggle to identify who is liable: the developer, the deploying institution, or the protocol itself. The subject of responsibility is absent, displaced into the codex. In medical auditing, compliance systems flag anomalies. Physicians defer to the system's

authority, administrators cite efficiency, and vendors disclaim liability by pointing to correct software operation. Patients face a null subject, unable to locate an accountable agent.

What these examples reveal is a shift from responsibility as **interpretive accountability** to responsibility as **structural absence**. In classical governance, even the most bureaucratic order carried a signature, a name, or an institutional office. Responsibility could be traced and contested. In executable governance, by contrast, responsibility is displaced at the very moment of execution. The compiled rule delivers authority while simultaneously erasing the accountable subject.

This displacement challenges legal traditions. Appeals presume a responsible subject who can review or overturn a decision. In predictive systems, appeals often reveal no subject to address. Courts can order explanations, but statistical opacity and syntactic delegation make explanations partial at best. Regulators can require documentation, but institutional shielding ensures that documentation confirms system-level compliance, not case-level responsibility. In each case, the structure of executable governance pushes responsibility outside the human circuit.

The implication is clear: responsibility will not reappear spontaneously within predictive systems. It must be deliberately reintroduced through accountability injection. Only by embedding human, hybrid, and syntactic-supervised tiers into the rule itself can responsibility be restored. Otherwise, displacement remains the structural default of executable governance.

4. Designing Accountability: The Three-Level Model

If the accountability gap is generated by syntactic delegation, statistical opacity, and institutional shielding, then responsibility cannot be restored by superficial oversight measures. It requires a structural redesign. This redesign is what I call *accountability injection*: the embedding of responsibility as a property of the *regla compilada* itself, equivalent to incorporating accountability into the generative grammar of governance.

Rather than treating responsibility as an external safeguard or post hoc review, accountability injection ensures that every decision produced by an executable system passes through a tiered framework where responsibility is formally reattached to human or institutional actors. This section introduces and develops the three-level model, human, hybrid, and syntactic supervised, showing how each level functions, when it applies, and how it prevents the displacement of responsibility.

4.1 Human Tier: Non Delegable Decisions

The human tier applies where outcomes are irreversible and rights bearing. These include criminal sanctions, life critical medical diagnoses, and exclusion from essential services such as housing or education. In such contexts, no amount of efficiency justifies the removal of human responsibility. A predictive system may propose, calculate, or flag, but the final decision must be tied to a named human subject who provides reasons.

Consider criminal sentencing. Predictive risk assessment tools, widely used in the United States, generate scores that suggest the likelihood of recidivism (Angwin, Larson, Mattu, & Kirchner, 2016, p. 4). Courts that rely mechanically on these scores risk delegating punitive authority to opaque algorithms. Under the human tier, a judge cannot simply adopt the model's output. They must document reasons for the decision, confirm the score's relevance, and sign their name. The signature is not symbolic; it re anchors responsibility in a subject. The predictive system supports, but never substitutes, human agency.

In healthcare, automated systems for radiological diagnostics can highlight suspicious anomalies. Yet the decision to diagnose cancer, initiate treatment, or deny care cannot be automated. A physician must make the determination, record their reasoning, and stand accountable. This prevents the patient from facing a null subject. If harm occurs, there is an identifiable actor to contest. This design principle embodies what the AI Act terms human oversight, but operationalized as non delegable decision rights (European Parliament, 2024, art. 14).

4.2 Hybrid Tier: Co Decision with Structural Safeguards

The hybrid tier addresses contexts where predictive efficiency is valuable but decisions remain materially impactful and potentially reversible. Examples include credit scoring,

insurance underwriting, and university admissions. Here the model can propose outcomes, but humans validate, dissent, or escalate.

A clear illustration is university admissions. Suppose an AI model ranks applicants by predicted academic success. In a purely executable regime, rejection letters are automatically generated. In the hybrid tier, however, the model produces both a recommendation and counterfactual scenarios, showing how rankings would differ if certain variables were weighted differently. Admissions officers then review this output, accept or override it, and record their decision. If they override, a structured dissent form is completed, explaining why. This dissent becomes part of the institutional record, ensuring traceability. The applicant now has an appeal path tied to identifiable actors rather than a statistical void.

In finance, credit scoring operates similarly. A predictive model may recommend lowering an applicant's credit limit. Before implementation, a human reviewer must validate the decision if the impact exceeds a defined threshold. If approved, the reviewer signs off; if not, the reviewer records dissent. This log becomes auditable, preventing institutions from hiding behind algorithmic opacity. Responsibility is not lost but layered, with both the system and the human reviewer leaving an accountable trace.

4.3 Syntactic Supervised Tier: Delegation with Traceability

The syntactic supervised tier applies to low risk or routine operations where delegation to a predictive system is acceptable, but only under conditions of strict traceability and automatic escalation triggers. Examples include fraud detection in small transactions, routine compliance audits, or document classification for archiving.

In this tier, decisions can be automated, but every execution is logged in an immutable ledger. The ledger must include inputs, model version, parameters, outputs, and any post processing steps. If anomalies are detected or if thresholds are crossed, the decision escalates automatically to the hybrid or human tier. For instance, in medical auditing, a system might automatically verify coding compliance for routine records. If the system flags a possible clinical risk, the decision escalates upward. In financial monitoring, small

irregularities may be processed syntactically, but larger or repeated anomalies trigger hybrid review.

This structure prevents responsibility from evaporating. Even where delegation is permitted, traceability guarantees that a responsible actor can be identified. Courts, auditors, and regulators can reconstruct the decision chain and attribute accountability. Responsibility becomes a compiled attribute of the decision process, not an optional addition.

5. Case Studies in Applied Governance

The theoretical foundation of accountability injection acquires meaning only when tested against real institutional contexts. This section examines four domains where executable governance is already present: the European Union's AI Act and its institutional architecture, smart contracts and decentralized autonomous organizations (DAOs), predictive scoring in education and finance, and automated medical audits. Each case illustrates both the displacement of responsibility and how the three-tier accountability model can reintroduce responsibility without dismantling efficiency.

5.1 AI Act and Institutional Architecture

The European Union's AI Act (Regulation EU 2024/1689) is the first comprehensive legislative attempt to regulate high-risk artificial intelligence systems. It establishes obligations for providers and deployers, ranging from risk management to human oversight (European Parliament, 2024, arts. 14, 28–30). Yet the Act operates at the level of conformity assessments and institutional designation, rather than at the point of decision. For example, Article 28 requires providers to ensure ongoing post-market monitoring, and Article 30 describes obligations of notifying authorities. These provisions guarantee that systems are assessed and documented, but they do not ensure that a patient denied care or a student rejected from admission can identify the responsible subject. The oversight remains diffuse.

Under the accountability injection model, the Act could be operationalized differently. Non-delegable decisions, such as those affecting fundamental rights, would fall under the human tier. Predictive admissions rankings would be processed in the hybrid tier, with structured dissent logs. Automated compliance checks could remain in the syntactic-supervised tier, but with immutable ledgers ensuring traceability. By layering accountability into decision processes, the Act would not merely regulate providers but also preserve appeal and remedy for individuals.

5.2 Smart Contracts and DAOs

Smart contracts exemplify the principle that “code is law” (Lessig, 1999, p. 6). They execute obligations automatically when predefined conditions are met. DAOs extend this logic by creating entire governance structures executed on-chain. Yet when disputes arise, responsibility is absent. A vulnerability exploited in 2016 drained millions of dollars from the Ethereum DAO. Courts and regulators struggled to determine liability: was it the developers, the participants, or the immutable code itself? Responsibility was displaced into syntax.

Accountability injection provides an alternative. In the human tier, high-stakes actions such as liquidating treasury funds would require named human authorization. In the hybrid tier, decisions like adjusting membership rules could be proposed algorithmically but approved by quorum with dissent logs. In the syntactic-supervised tier, routine payouts or fee adjustments could remain automatic, but immutable logs and escalation rules would allow disputes to be traced back to responsible actors. Instead of erasing accountability, the DAO would preserve efficiency while anchoring responsibility.

5.3 Predictive Scoring in Education and Finance

University admissions systems increasingly rely on predictive models to rank applicants. These systems promise efficiency but generate opaque and unappealable outcomes. Applicants denied admission often receive generic explanations. Under the hybrid tier, the model would propose rankings, but admissions officers would validate them and record dissent when applicable. A rejected student could then appeal to an identifiable actor who approved the decision.

In finance, credit scoring systems can reduce or deny credit limits based on opaque correlations. Here too, the hybrid tier ensures validation by a responsible human. If the model recommends a significant change, a reviewer must approve and document the decision. Small routine adjustments may remain under the syntactic-supervised tier, but escalations trigger human review. This preserves both efficiency and accountability.

5.4 Automated Medical Audits

Hospitals and insurers increasingly employ automated auditing systems to flag anomalies in billing and diagnostics. These systems reduce costs but risk false positives and denials. When patients appeal, institutions often point to the system, vendors deny liability, and regulators confirm compliance. Responsibility is diffused.

The three-tier model addresses this. Low-risk audits, such as coding checks for routine procedures, remain in the syntactic-supervised tier with full traceability. Higher-risk anomalies escalate to the hybrid tier, where medical administrators validate or dissent. Life-critical findings automatically escalate to the human tier, requiring a physician's signed decision. The patient gains a clear appeal path. Institutions maintain efficiency while reintroducing accountability.

5.5 Comparative Insight

Across all four cases, the same pattern emerges. Predictive systems displace responsibility by making structure sufficient for authority. The AI Act enforces oversight but without decision-level accountability. DAOs enforce contracts automatically, erasing subjects. Predictive scoring delivers efficiency at the cost of appeal. Automated medical audits create binding outcomes without responsible agents. The accountability injection model demonstrates how non-delegable, hybrid, and syntactic-supervised tiers can restore responsibility without disabling efficiency.

By embedding these tiers into the *regla compilada* of governance systems, institutions avoid the trap of spectral sovereignty and null subjects. Accountability becomes a structural property rather than an external safeguard. Individuals regain the right to appeal.

Institutions regain legitimacy. Predictive societies can remain efficient while remaining responsible.

6. Ethical and Legal Implications

The introduction of accountability injection through the three-level model transforms the landscape of predictive governance, but it also raises profound ethical and legal implications. This section examines how the framework reintroduces appeal and responsibility, evaluates its compatibility with existing legal standards, and considers the broader consequences for institutional legitimacy in predictive societies. The argument is that accountability must be conceived not as an external correction but as a compiled property of governance systems themselves.

6.1 Reintroducing Appeal and Pragmatic Responsibility

One of the defining features of classical legal systems is the possibility of appeal. An administrative order, a judicial sentence, or a regulatory decision can be contested before another authority. The right to appeal is essential not only for justice but also for legitimacy, since it assures affected individuals that no decision is beyond contestation (Kelsen, 1967, p. 114). Predictive governance erodes this mechanism. When a decision is produced by a model or compiled code, there is often no subject to whom appeal can be directed. The three-level model restores appeal by binding decisions to responsible tiers. In the human tier, the responsible actor is explicit. In the hybrid tier, dissent logs provide a documented path of contestation. In the syntactic-supervised tier, immutable ledgers allow individuals to reconstruct the decision chain. Appeal thus re-enters execution, not as an external safeguard but as a structural property of the decision-making process.

For example, in credit scoring, individuals often receive opaque outcomes with no clear path of appeal. Under accountability injection, if a decision is validated by a human reviewer in the hybrid tier, the reviewer's signature and reasoning provide a concrete point of contestation. In medical audits, if an automated system flags an anomaly, escalation rules ensure that a physician or administrator signs off on high-risk outcomes. Patients can

therefore appeal directly to a responsible actor, rather than facing the void of a null subject. This reintroduction of appeal transforms predictive governance into a space where responsibility is not displaced but redistributed in formal tiers.

6.2 Compatibility with Legal Standards and International Norms

Embedding accountability into executable governance aligns with and extends current legal frameworks. The European Union’s AI Act requires human oversight for high-risk systems (European Parliament, 2024, art. 14). However, oversight is vaguely defined, and conformity assessments focus on system-level compliance rather than case-level responsibility. The three-level model operationalizes oversight by clarifying when decisions must remain human, when hybrid validation applies, and when supervised delegation is permissible. In this sense, accountability injection serves as a normative blueprint that could inform future regulatory guidance and jurisprudence.

Beyond the EU, international organizations such as UNESCO and the OECD have emphasized transparency and accountability in AI ethics (OECD, 2019, p. 12). Yet their principles often remain aspirational. The three-level model provides a concrete mechanism for operationalizing these principles. For instance, OECD guidelines stress the need for human-centric AI, but without specifying how to embed human judgment structurally. Accountability injection answers this by codifying human, hybrid, and syntactic-supervised tiers directly into the execution pipeline. This makes accountability auditable, reproducible, and legally enforceable.

6.3 Risks and Trade-offs

Embedding accountability raises risks and trade-offs. One concern is efficiency loss. Human and hybrid tiers may slow down decision-making, introducing delays in contexts where speed is valued. In financial trading, for example, requiring human validation for every transaction could paralyze operations. To mitigate this, the model restricts human tiers to non-delegable, high-stakes decisions, while allowing syntactic supervision for low-risk operations. Another concern is over-escalation. If escalation thresholds are too sensitive, many decisions may default to the human tier, overwhelming institutions. This can be mitigated by calibrating thresholds empirically and reviewing them periodically.

There is also the risk of symbolic compliance, where institutions implement formal logs without genuine accountability. To prevent this, immutable ledgers and structured dissent must be auditable by external regulators. If dissent forms are systematically ignored, or if ledgers are incomplete, institutions can be held legally liable. In this sense, accountability injection introduces enforceable duties rather than voluntary practices.

6.4 Toward International Accountability Standards

The broader implication is the possibility of developing international accountability standards for predictive governance. Just as the International Organization for Standardization (ISO) provides benchmarks for quality management, accountability injection could serve as the basis for certifiable standards. A system could be audited for compliance with tier allocation, escalation rules, and decision ledgers. This would harmonize accountability across jurisdictions, ensuring that predictive governance remains legitimate even in transnational contexts such as finance, healthcare, or education.

By integrating accountability into the *regla compilada*, the model also advances ethical reflection. It shifts responsibility from being reactive, a matter of liability after harm, to being proactive, a matter of structural design. In predictive societies, where authority is increasingly produced by form rather than by deliberation, ethics cannot remain external. Ethics must be compiled, just as rules are compiled. Accountability injection is therefore not merely a legal innovation but an ethical imperative.

7. Conclusion: Restoring Responsibility in Predictive Societies

The analysis presented throughout this article demonstrates that executable governance has created a structural accountability gap. Authority is increasingly produced by syntax, embedded in predictive models and compiled rules, while the responsible subject disappears. This produces what has been called *spectral sovereignty* and *null subjects of power* (Startari, 2025a; Startari, 2025b). Decisions bind institutions and individuals, yet no actor can be identified or appealed to. Traditional oversight frameworks, such as those offered by the AI Act, attempt to address this through institutional monitoring and

conformity assessments (European Parliament, 2024). However, they often fail to secure responsibility at the level of specific decisions. The conclusion of this article is that only through accountability injection, compiled directly into the *regla compilada*, can responsibility be structurally restored in predictive societies.

7.1 Formal Definition of Accountability Injection

Accountability injection refers to the embedding of responsibility into executable governance systems as a compiled attribute rather than an external safeguard. It ensures that each decision produced by a predictive system is structurally tied to a responsible subject through one of three tiers. The human tier secures non-delegable decisions, the hybrid tier creates co-decision processes with traceable dissent, and the syntactic-supervised tier permits delegation under strict conditions of traceability and escalation. Together these tiers guarantee that no decision is executed without the possibility of appeal and attribution. Responsibility is no longer presumed to follow authority; it is compiled as a formal property of governance.

7.2 Reconstructing the Link Between Legitimacy and Responsibility

From Weber's rational-legal authority to Kelsen's hierarchy of norms, legitimacy has historically been tied to responsibility (Weber, 1978; Kelsen, 1967). Predictive governance disrupts this link, producing legitimacy without accountability. The model developed here reconstructs the bond by embedding accountability directly into the generative grammar of governance. This restores the classical structure: authority is legitimate only if it is accountable. By compiling accountability into the rule, predictive governance regains the possibility of legitimacy without sacrificing efficiency.

7.3 Implications for Legal and Institutional Design

For legislators and regulators, the implication is that laws must do more than mandate oversight in general terms. They must prescribe accountability injection mechanisms at the level of decision processes. This could mean requiring dissent logs in admissions systems, physician sign-off in medical audits, or immutable ledgers in financial compliance. For courts, the model provides a framework for attributing responsibility even when decisions are produced by code. A complete decision ledger can be audited to determine who

validated or overrode a decision. For institutions, accountability injection offers a way to regain legitimacy in the eyes of citizens who face opaque, automated decisions.

7.4 Ethical Imperatives in Predictive Societies

The ethical implications are equally significant. In societies where authority is produced by predictive systems, ethics cannot be external. Appeals to fairness or transparency are insufficient if they are not embedded in the decision-making process itself. Accountability injection shifts ethics from an aspirational layer to a compiled feature. It ensures that the dignity of individuals is respected by guaranteeing appeal and responsibility. In this sense, accountability injection represents an ethical imperative as much as a legal one.

7.5 Future Integration of Syntactic Ethics

Looking forward, the model invites integration with international standards. Just as ISO norms provide benchmarks for quality management, accountability injection could form the basis of certifiable standards for AI governance. Regulators could require certification of systems according to their compliance with tier allocation, escalation rules, and ledger integrity. This would harmonize accountability practices across jurisdictions. Moreover, the framework opens a new field of *syntactic ethics*, where ethical obligations are formalized and compiled into governance systems. Ethics becomes a property of structure, not an external discourse.

7.6 Closing Reflection

The central claim of this article is falsifiable. If appeals and escalation mechanisms are compiled into executable governance, institutions regain identifiable responsible subjects. If they are absent, responsibility dissolves. In this sense, accountability injection is not only a theoretical construct but also an operational test. Its adoption can be verified through documentation, ledgers, and appeal processes. Predictive societies stand at a threshold. Without accountability injection, they risk entrenching spectral sovereignty and null subjects. With it, they can restore responsibility and legitimacy while maintaining efficiency.

By embedding accountability directly into the *regla compilada*, predictive governance can move beyond opacity and irresponsibility. It can achieve a synthesis where efficiency and

responsibility co-exist. The restoration of responsibility is not nostalgic, seeking to return to pre-digital forms of governance. It is structural and forward-looking, offering a blueprint for institutions that must govern in an age where authority is generated by form. In this sense, accountability injection is not only a legal and ethical innovation but also a necessary condition for preserving legitimacy in predictive societies.

References

Angwin, J., Larson, J., Mattu, S., & Kirchner, L. (2016). Machine bias: There is software used across the country to predict future criminals. And it is biased against blacks. *ProPublica*. <https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing>

Burrell, J. (2016). How the machine “thinks”: Understanding opacity in machine learning algorithms. *Big Data & Society*, 3(1), 1–12. <https://doi.org/10.1177/2053951715622512>

Chomsky, N. (1965). *Aspects of the theory of syntax*. Cambridge, MA: MIT Press.

European Parliament. (2024). Regulation (EU) 2024/1689 of the European Parliament and of the Council of 12 July 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). *Official Journal of the European Union*, L, 2024, 1–127.

Kelsen, H. (1967). *Pure theory of law* (2nd ed.). Berkeley: University of California Press.

Lessig, L. (1999). *Code and other laws of cyberspace*. New York: Basic Books.

Montague, R. (1974). *Formal philosophy: Selected papers of Richard Montague*. New Haven: Yale University Press.

OECD. (2019). *OECD principles on artificial intelligence*. Paris: OECD Publishing. <https://doi.org/10.1787/eedfee77-en>

Startari, A. V. (2025a). *Spectral Sovereignty: Authority Without Presence in Predictive Systems*. SSRN Electronic Journal. <https://doi.org/10.5281/zenodo.17026629>

Startari, A. V. (2025b). *Null Subjects of Power: Disappearing Responsibility in Executable Governance*. SSRN Electronic Journal. <https://doi.org/10.5281/zenodo.17085900>

Startari, A. V. (2025c). *The Codex of Authority: Legitimacy Through Syntactic Form*. SSRN Electronic Journal. <https://doi.org/10.5281/zenodo.17026629>

Startari, A. V. (2025d). *Ethos Without Source: Algorithmic Identity and the Simulation of Credibility*. SSRN Electronic Journal. <https://doi.org/10.2139/ssrn.5313317>

Weber, M. (1978). *Economy and society* (G. Roth & C. Wittich, Eds.). Berkeley: University of California Press.

