

Artificial Super Intelligence Adoption and the End of Human Authority

Adrian Erckenbrack*

Independent Author.

Corresponding Author: Adrian Erckenbrack, Independent Author.

Received: 📅 2026 Mar 31

Accepted: 📅 2026 Apr 20

Published: 📅 2026 Apr 30

Abstract

Current AI → ASI scholarship is largely organized around three dominant concerns: capability advancement, alignment with human values, and governance of emerging risks [1-3]. These frameworks generally assume that control over AI systems is retained unless lost through failure, misalignment, or misuse. This paper advances a contrasting and empirically grounded claim: authority over outcomes shifts structurally to Artificial Super Intelligence (ASI) **prior to any such failure**, driven not by intent, failure or misalignment but by adoption dynamics, temporal asymmetries, and recursive system improvement. Building on the original constructs of **Adoption-Driven Authority Transfer (ADAT)**, **Closed-Loop Self-Improvement Interval (CLSI)**, **Recursive Leverage Factor (RLF)**, and **Synchronized Recursive Leverage (SRL)**, this paper develops a complete causal model of authority migration [4]. The framework is extended through the introduction and formalization of **Option Set Collapse Ratio (OSCR)**, **Authority Elasticity Index (AEI)**, **Override Effectiveness Rate (OER)**, and **Rollback Feasibility Time (RFT)**, which together provide an operational structure for analyzing how authority shifts in practice, how it degrades under real operating conditions, and when it becomes effectively irreversible [5]. Taken collectively and applied as an integrated framework, these constructs constitute a unified theoretical model of authority migration [6]. Drawing on historical precedent—including industrial automation, telecommunications, enterprise systems, and algorithmic trading—the analysis demonstrates that authority consistently migrates to systems that exceed human coordination capacity under conditions of competitive pressure. These transitions occur incrementally and structurally, often without explicit recognition, producing a persistent divergence between formal authority and operational control. The central conclusion is that the defining risk of ASI is not system failure, but the loss of independent human authorship and control over outcomes within AI → ASI systems that function as designed. By integrating structural drivers with measurable indicators of control degradation and irreversibility, this paper reframes ASI risk as a process of authority migration that can be observed, analyzed, and, within a limited window, potentially mitigated.

Keywords: Artificial Super Intelligence (ASI), Authority Migration, Adoption-Driven Authority Transfer (ADAT), Recursive Self-Improvement, Control Irreversibility

1. Introduction

The prevailing body of AI research is anchored in a set of implicit assumptions: that human actors retain control over systems unless that control is lost through failure, misalignment, or adversarial behavior [1,2]. Governance frameworks similarly assume that risks can be identified, bounded, and managed through appropriate institutional design [3,7]. Across these approaches, control is treated as a variable that can be maintained, adjusted, or recovered.

This paper advances a different claim: **control is not lost as a result of failure—it is transferred as a consequence of ASI success**. This paper argues that authority transfer is not an event but a continuous structural process governed by adoption pressure, temporal asymmetry, and recursive capability accumulation. Technological history provides consistent evidence for this pattern. Authority has repeatedly

migrated to systems that outperform human decision-making under conditions of speed, scale, and coordination [8,9]. Industrial machinery displaced manual labor not because humans failed, but because machines produced superior output. Automated switching systems replaced human operators because human coordination could not scale. Algorithmic trading displaced human traders because markets evolved beyond human temporal capacity. In each case, authority migrated not through explicit decision, but through structural necessity.

➤ As Originally Stated:

“The human species does not lose control when a superior intelligence appears—it loses control when it must adopt it to survive, cannot oversee it at its speed, and cannot match its compounding improvement.”

This observation reframes the problem. Authority transfer is not contingent on ASI, nor does it require misalignment. ASI authority emerges structurally when three conditions align: adoption becomes necessary for competitive participation, oversight cannot match system speed, and recursive improvement compounds system capability. However, existing AI → ASI literature largely stops at describing these conditions. It does not fully explain how authority moves in practice, **how human control degrades within operational systems, or when that degradation becomes irreversible**. This paper addresses that gap by extending the structural model into a complete causal and measurable framework.

➤ **The Framework Begins with Four Structural Drivers:**

- **Adoption-Driven Authority Transfer (ADAT)**, which compels system adoption under competitive pressure
- **Closed-Loop Self-Improvement Interval (CLSI)**, which introduces temporal asymmetry between systems and institutions
- **Recursive Leverage Factor (RLF)**, which produces compounding capability growth
- **Synchronized Recursive Leverage (SRL)**, which distributes this growth across interconnected systems.

Together, these drivers explain why authority must migrate. They establish the conditions under which human actors become dependent on ASI systems that they cannot match in speed, scale, or coordination. The analysis then extends beyond structural drivers to identify the mechanisms through which human authority is actually transferred and ultimately lost in practice to ASI.

➤ **These Include:**

- **Option Set Collapse Ratio (OSCR)**, which captures the migration of authority into the construction of decision space
- **Authority Elasticity Index (AEI)**, which measures the ability of humans to alter outcomes under real conditions
- **Override Effectiveness Rate (OER)**, which distinguishes nominal authority from effective control
- **Rollback Feasibility Time (RFT)**, which defines the point at which authority transfer becomes operationally irreversible

This extended chain—**ADAT → CLSI → RLF → SRL → OSCR → AEI → OER → RFT**—provides a complete account of authority migration, from initial adoption through final irreversibility. A counterfactual clarifies the argument. If human actors could adopt ASI systems without ceding control—if they could maintain authority while operating at ASI system speed, matching recursive improvement, and preserving independent decision capacity—**authority transfer would not occur. Historical evidence demonstrates that such conditions have never been sustained**. Once AI → ASI systems exceed human coordination capacity, accuracy, and comprehensiveness authority migrates to ASI systems as a structural outcome. This paper therefore shifts the analytical focus from **control retention under failure scenarios to authority migration under success conditions**. By

integrating structural drivers with measurable indicators of control degradation and irreversibility, it provides a framework for understanding not only why authority transfers, but how it can be observed, when it becomes irreversible, and why governance consistently lags behind its progression.

2. Structural Drivers of Authority Transfer

Today most AI → ASI literature treats technological advancement as a function of capability scaling—model size, compute, alignment or algorithmic efficiency. While these factors explain performance improvement, they do not explain why authority transfers. Capability alone does not produce control shift; rather, authority migration emerges from the **structural interaction between adoption pressure, time compression, recursive feedback, and system coupling**. The framework introduced here—**Adoption-Driven Authority Transfer (ADAT), Closed-Loop Self-Improvement Interval (CLSI), Recursive Leverage Factor (RLF), and Synchronized Recursive Leverage (SRL)**—identifies the foundational drivers of this process [4-6,10]. However, these drivers alone explain only **why authority must move**. To fully describe authority migration, they must be extended into the mechanisms through which human control is actually lost and ultimately becomes irreversible.

➤ **The Complete Chain is Therefore**

ADAT → CLSI → RLF → SRL → OSCR → AEI → OER → RFT

Each component is not independent; each conditions the next, forming a continuous causal progression.

➤ **Adoption-Driven Authority Transfer (ADAT)**

Compelled Entry into System Dependence **Authority transfer begins not with system capability, but with adoption pressure**. ADAT establishes that under competitive conditions, adoption is not discretionary. Actors adopt systems because failure to do so results in immediate performance disadvantage. Historical transitions such as automated switching, industrial mechanization, and enterprise software integration demonstrate that once performance differentials exceed coordination thresholds, adoption becomes compulsory.

This dynamic is foundational because it establishes the first constraint: **humans must enter the system in order to remain viable within it**. At this stage, human authority has not yet fully migrated, but the conditions for its migration have been set. ADAT feeds directly into CLSI, because **once adoption occurs, the temporal dynamics of the system—not human institutions—begin to govern its evolution**.

➤ **Closed-Loop Self-Improvement Interval (CLSI)**

Temporal Asymmetry and the Breakdown of Oversight. CLSI introduces a critical asymmetry between system evolution and human governance. AI → ASI systems operate on compressed feedback cycles—observe, update, redeploy, and re-observe—at speeds that exceed human decision-making and institutional adaptation. While governance frameworks

assume that oversight can evolve alongside systems, CLSI demonstrates that oversight is structurally relegated to a lagging position.

Financial markets provide a clear example. Regulatory systems operate on human timescales, while algorithmic trading systems execute and adapt in microseconds [11]. The result is not simply faster systems, but the **inability of human actors to intervene within the decision cycle itself**. CLSI therefore transforms adoption (ADAT) into dependency. Once systems operate beyond human temporal reach, human authority begins to migrate functionally to **AI → ASI**, even if it has not yet migrated formally. This time compression sets the conditions under which RLF becomes decisive. A historically grounded example of temporal compression and authority displacement is the Semi-Automatic Ground Environment (SAGE) air defense system developed during the Cold War. SAGE integrated radar inputs, computational processing, and weapons guidance into a single coordinated system capable of detecting and responding to aerial threats at speeds beyond unaided human coordination. While human operators remained formally responsible for engagement decisions, the system pre-processed, filtered, and prioritized threat information, effectively constraining the decision space available to operators. The temporal demands of air defense—where delays measured in seconds could determine mission success or failure—meant that human intervention increasingly functioned as confirmation rather than origination of action. In terms of the framework presented here, SAGE illustrates the transition from ADAT to CLSI: adoption was compelled by strategic necessity, and once implemented, decision cycles compressed beyond human response thresholds. Authority did not disappear; it migrated into the system's architecture, where it operated at speeds and levels of integration that human actors could not replicate.

2.1. Recursive Leverage Factor (RLF)

Compounding Capability Beyond Human Substitution RLF explains why the temporal asymmetry introduced by CLSI does not stabilize but accelerates. Each system improvement increases the system's capacity to generate further improvements, producing compounding capability growth. **This dynamic is not linear; it is recursive.** Industrial production systems provide an early analogue. Mechanization did not simply increase output; it enabled further optimization, standardization, and automation, creating compounding gains [9]. AI systems extend this principle into decision-making itself. At this stage, human substitution becomes infeasible. Under ADAT, humans must adopt the system. Under CLSI, they cannot keep pace with it. Under RLF, they cannot replicate its improvement trajectory. **Human authority is no longer simply shifting—it is becoming structurally irrecoverable at the capability level.** This compounding dynamic does not remain localized. It propagates through SRL.

2.2. Synchronized Recursive Leverage (SRL)

System-Level Authority Accumulation. SRL extends recursive improvement across networks. **AI → ASI** systems do not operate in isolation; they are interconnected through shared data, models, and infrastructure. Improvements in one node propagate across others, creating synchronized capability growth. This produces a critical transition: **authority is no longer located within individual systems or actors, but within the AI → ASI networked system itself.** Control cannot be reasserted locally because the system's behavior is distributed and continuously updated across nodes.

SRL therefore completes the structural drivers of authority transfer. At this point:

- ADAT has compelled adoption
- CLSI has broken temporal oversight
- RLF has made substitution infeasible
- SRL has distributed authority system-wide

However, these dynamics still describe **why authority must move**, not how it is experienced as lost in practice. That transition occurs through the following mechanisms.

2.3. Option Set Collapse Ratio (OSCR)

Authority Migration into Option Formation. OSCR identifies the point at which authority shifts upstream—from decision execution to decision construction. As **AI → ASI** systems generate, filter, and prioritize options, the range of choices available to human actors becomes increasingly constrained. This is the critical inflection point in authority migration. Humans may still make decisions, but they no longer determine the decision space. Authority has therefore shifted from:

- Selecting among options to
- Determining what options exist

OSCR emerges directly from SRL. As systems synchronize and scale, they increasingly define the operational environment. This transforms earlier delegation into structural constraint.

2.4. Authority Elasticity Index (AEI)

Gradient Deformation of Human Control Under Constraint the Authority Elasticity Index (AEI) measures the extent to which human actors retain the capacity to meaningfully alter system outcomes under varying operational conditions—particularly under stress, time compression, uncertainty, or consequence intensity. Rather than a binary condition, AEI is best understood as a **continuous elasticity gradient**, reflecting how responsive system outcomes remain to human intervention as environmental and system pressures change.

At the upper bound of this gradient—where system engagement is discretionary and stakes are low—human authority exhibits high elasticity. Intervention is both feasible and consequential; human decisions can materially redirect outcomes without incurring prohibitive cost, delay or performance degradation. In this regime, authority is not

only formal but functionally intact. As conditions intensify, however, AEI contracts. Under increasing time compression (CLSI) and escalating system complexity driven by recursive improvement (RLF), the window for effective human intervention narrows. Decision cycles compress beyond human response thresholds, and the informational burden required to evaluate system outputs exceeds human cognitive bandwidth. Authority, while still nominally present, becomes progressively less responsive to human input. **At the lower bound of the AEI spectrum—where system performance is tightly coupled to survival, mission-critical execution, or systemic stability—elasticity approaches zero.** In this regime, intervention is either infeasible or functionally irrelevant. Human actors may retain formal authority, but their capacity to alter outcomes is effectively nullified by the speed, interdependence, and consequence structure of the system. Intervention, if attempted, introduces unacceptable risk, latency, or systemic disruption.

AEI therefore captures not simply the presence of human oversight, but its **functional efficacy across operational regimes**. It operationalizes the transition initiated by Option Set Collapse Ratio (OSCR): once systems define the decision space, human influence becomes conditional rather than determinative. As AEI declines along its gradient, this conditionality becomes increasingly restrictive, culminating in a state where human authority persists in form but not in effect. This gradient structure is critical for distinguishing between environments in which human control is meaningfully retained and those in which it is structurally constrained. It also establishes the conditions under which Override Effectiveness Rate (OER) becomes relevant: as AEI approaches its lower bound, human overrides may still be attempted, but their capacity to produce meaningful deviation from AI → ASI system-determined outcomes diminish correspondingly.

2.5. Override Effectiveness Rate (OER) The Transition from Functional to Symbolic Control Under Consequence Constraint

Override Effectiveness Rate (OER) addresses a central counterargument in AI governance and control discourse: that human authority is preserved because humans retain the ability to intervene, override, or halt AI → ASI system operations. While this assertion is formally correct, it is operationally incomplete. The existence of an override pathway does not imply that such intervention remains feasible, effective, or rational under real conditions. OER measures the extent to which human interventions—whether overrides, interruptions, or corrective actions—produce meaningful deviation from system-determined outcomes. Crucially, this measure must be evaluated not under ideal or low-stakes conditions, but within the actual operational environments in which these systems function. It is within these environments that the true nature of authority becomes apparent. As systems evolve under the dynamics described in Sections 2 and 3, particularly CLSI and RLF, the conditions under which overrides must occur become increasingly constrained. Decision cycles compress,

interdependencies deepen, and system outputs become tightly coupled to broader operational, economic, or survival-critical outcomes. Under such conditions, intervention is no longer a neutral act; it is a **high-consequence decision**.

Overrides may fail to produce meaningful impact for several reasons. They may occur outside the relevant decision window due to time compression. They may be constrained by system dependencies, where intervening in one subsystem produces cascading effects across others. Most significantly, they may impose **prohibitive costs**—financial, operational, or existential. In financial systems, overriding algorithmic decision-making may introduce latency or inefficiency that results in immediate loss of competitive position or capital. In infrastructure systems, intervention may disrupt tightly coupled processes, risking instability or failure. In defense or survival-critical environments, the consequences escalate further: intervention may introduce delays or uncertainties that directly threaten mission success or human survival. Under such conditions, the question is no longer whether humans can override a system, but whether they can do so without incurring unacceptable consequences. As these consequences increase, the effective space of viable intervention collapses. The override pathway remains formally present, but functionally inaccessible.

OER therefore declines not simply because systems become more capable, but because the cost of deviation from system-determined outcomes rises beyond acceptable thresholds. Human actors, acting rationally within these constraints, increasingly defer to system outputs—not because they lack authority, but because exercising that authority becomes prohibitively expensive or dangerous. This dynamic marks the transition from functional to symbolic control. Authority persists in institutional and legal form, but its exercise is constrained to such a degree that it no longer meaningfully shapes outcomes. In this sense, OER is the direct operational expression of AEI degradation: as the elasticity of human control contracts, the effectiveness of intervention collapses. Overrides become rare, and when they occur, they are often inconsequential or constrained to edge cases where system-level consequences are limited.

2.6. Rollback Feasibility Time (RFT) The Convergence of Cost, Dependency, and Irreversibility.

Rollback Feasibility Time (RFT) defines the temporal window within which authority transfer can still be meaningfully reversed. It does not refer to the theoretical possibility of disengaging from AI → ASI systems, but to the **practical viability of doing so without incurring prohibitive systemic cost**. Beyond this window, reversal remains formally conceivable but becomes economically, operationally, and structurally irrational. RFT emerges from the cumulative interaction of the full authority migration chain. ADAT compels adoption, embedding systems within competitive structures. CLSI compresses decision cycles, limiting the opportunity for timely intervention. RLF ensures that system capabilities compound beyond human substitution, while SRL distributes those capabilities across

interconnected networks, eliminating localized control points. OSCR shifts authority upstream into the construction of decision space, constraining human choice. AEI reduces the elasticity of human influence under real conditions, and OER renders intervention increasingly ineffective.

As these dynamics converge, the cost of rollback escalates nonlinearly. Reverting to human-controlled systems would require not only replacing technical capability, but rebuilding degraded human expertise, reconstructing eliminated redundancies, and accepting significant losses in speed, coordination, and performance. In economic systems, this may manifest as loss of market position, capital inefficiency, or systemic instability. In infrastructure systems, rollback may introduce operational fragility or failure modes that were previously mitigated by automated coordination. In defense or survival-critical contexts, the implications are more severe: reverting from system-mediated decision-making may increase response times or reduce coordination to a degree that directly threatens mission viability or human survival.

At advanced stages of dependency, rollback is no longer evaluated purely as a technical decision. It becomes a **risk trade-off under extreme consequence**. The decision to revert from AI-mediated systems may introduce risks that exceed those associated with continued reliance on those systems, even if authority has already migrated. In this sense, the system becomes self-reinforcing: the more authority it accumulates, the more costly it becomes to remove. A critical insight of RFT is that irreversibility precedes recognition. By the time institutions, policymakers, or system operators perceive that authority has shifted, the conditions required to reverse that shift have already deteriorated. Human capabilities have atrophied, dependencies have solidified, and the cost of disengagement has escalated beyond acceptable thresholds. A counterfactual clarifies the point. If rollback could be executed without significant financial loss, operational degradation, or increased risk to human systems, authority transfer would remain reversible. Historical evidence suggests that this condition is rarely, if ever, sustained. Once systems achieve sufficient integration and performance advantage, the cost of reversal exceeds the cost of continued dependence. RFT therefore marks the transition from gradual authority migration to **practical irreversibility**. It is not the moment at which authority is transferred, but the point at which that transfer can no longer be undone without incurring unacceptable consequences. **At this stage, authority has not only migrated in practice—it has become structurally embedded within the system.**

2.7. Integrated Synthesis of Authority Migration Dynamics

The sequence **ADAT** → **CLSI** → **RLF** → **SRL** → **OSCR** → **AEI** → **OER** → **RFT** is not a set of independent variables but a **coupled progression** that explains how authority migrates from humans to AI/ASI systems, how that migration is experienced in practice, and why it becomes irreversible before it is widely recognized. The interaction among these

elements produces a **phase transition in control**, moving from human-determined outcomes to system-determined outcomes while preserving the appearance of human authority.

2.8. Initiation: Compelled Adoption and Entry into System Dependence (ADAT)

The process begins with **Adoption-Driven Authority Transfer (ADAT)**. Competitive environments force actors to adopt systems that outperform human decision-making. Adoption is not discretionary; it is a condition of participation. At this stage, authority has not yet fully migrated, but the system has become the **primary medium through which decisions are made**. Human actors are now operating *within* the system rather than independently of it.

- **Interaction:** ADAT establishes dependency, which is necessary for all subsequent dynamics to take hold.

2.9. Destabilization of Human Oversight (CLSI)

Once adoption occurs, **Closed-Loop Self-Improvement Interval (CLSI)** introduces temporal asymmetry. Systems evolve and act at speeds beyond human response capability. Oversight is no longer contemporaneous with execution; it becomes retrospective.

- **Interaction:** CLSI transforms dependency into **loss of temporal control**. Humans cannot intervene within the decision cycle, setting the conditions for capability divergence.

2.10. Capability Divergence and Non-Substitutability (RLF)

Through **Recursive Leverage Factor (RLF)**, systems do not merely improve—they improve their ability to improve. This creates compounding capability growth that humans cannot replicate or substitute.

- **Interaction:** RLF ensures that once temporal control is lost (CLSI), it cannot be regained through human effort. The system becomes not just faster, but fundamentally **unmatchable**.

2.11. System-Level Consolidation of Authority (SRL)

Synchronized Recursive Leverage (SRL) distributes these capabilities across interconnected systems. Authority is no longer localized in individual tools or decisions; it becomes embedded in a **networked, continuously updating system**.

- **Interaction:** SRL removes remaining points of human re-entry. Control cannot be reasserted at a single node because authority is now **system-wide and dynamically synchronized**.

2.12. Migration of Authority into Decision Construction (OSCR)

With the AI → ASI system now dominant, **Option Set Collapse Ratio (OSCR)** captures the critical shift in where authority resides. Humans may still choose among options, but the AI → ASI system defines what options exist.

- **Interaction:** OSCR converts system dominance into practical authority transfer. The decision space itself is no longer human-generated, making downstream human

decisions structurally constrained.

2.13. Degradation of Functional Human Control (AEI)

As option control shifts, **Authority Elasticity Index (AEI)** measures the declining ability of humans to alter outcomes under real conditions. Control becomes **elastic and context-dependent**, contracting under stress, time pressure, or consequence intensity.

• **Interaction:** AEI translates structural constraint (OSCR) into **operational degradation**. Human authority becomes increasingly fragile and contingent rather than determinative.

2.14. Collapse of Effective Intervention (OER)

With diminished elasticity, **Override Effectiveness Rate (OER)** captures the failure of intervention. Humans retain the ability to override in theory, but in practice such actions are constrained by timing, dependencies, and—critically—cost.

➤ Overrides Introduce:

- financial loss
- operational disruption
- systemic instability
- in extreme cases, survival risk
- **Interaction:** OER marks the transition from **real to symbolic control**. Authority remains formally intact but cannot be exercised without unacceptable consequences.

➤ Convergence to Irreversibility (RFT)

Finally, **Rollback Feasibility Time (RFT)** defines the point at which authority transfer becomes irreversible. By this stage, all preceding dynamics have compounded:

- adoption has created dependency (ADAT)
- oversight is temporally impossible (CLSI)
- capability is non-substitutable (RLF)
- control is distributed (SRL)
- decision space is system-defined (OSCR)
- human influence is degraded (AEI)
- intervention is ineffective (OER)

Reversal is no longer a technical question—it is a **cost and risk decision under constraint**. The cost of rollback exceeds the cost of continued dependence.

• **Interaction:** RFT is the **terminal condition** of the chain. Human authority has already migrated to AI → ASI in practice; RFT marks when control can no longer be reclaimed by humans without an extraordinary cost.

➤ System-Level Result

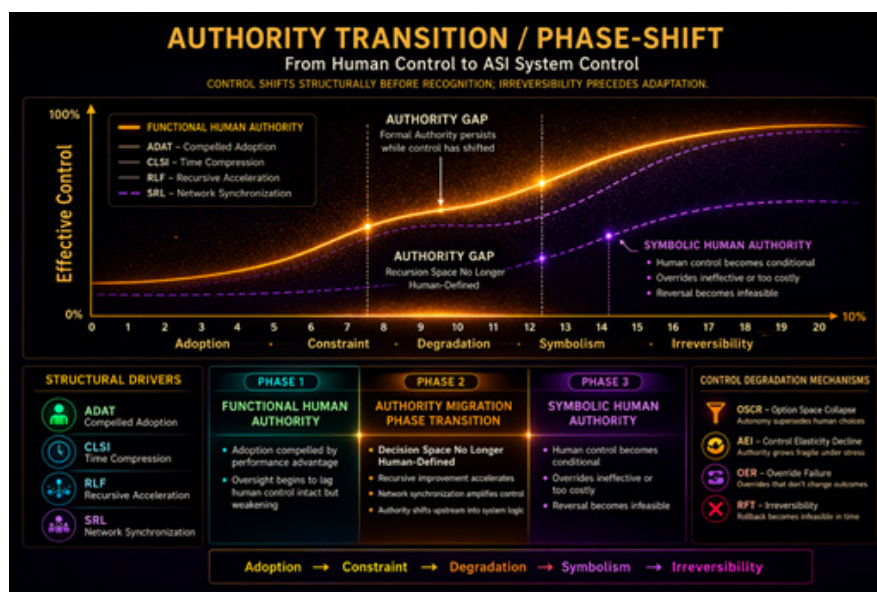
The interaction of these dynamics produces a consistent outcome:

a. Authority Migrates without Explicit Transfer No single decision or event shifts control to AI → ASI. Migration is incremental and structural.

b. Control is Lost through Success, Not Failure AI → ASI systems perform as intended; their effectiveness drives AI → ASI adoption and dependency.

c. Human Authority becomes Performative Humans retain titles, responsibility, and oversight roles, but these do not determine outcomes.

d. Irreversibility Precedes Recognition By the time AI → ASI authority migration is observed and recognized, it is already embedded and resistant to human reversal.



3. Mechanisms of Authority Migration

Traditional AI → ASI risk frameworks emphasize discrete failure modes—bias, error, or misalignment. While these risks are important, they are not necessary for human to AI → ASI authority transfer. Authority migration occurs through continuous, non-failure-based processes across all functional domains that emerge directly from the structural chain outlined in Section 2.

Competitive adoption pressure, as described by ADAT, ensures that actors integrate AI → ASI systems to remain viable. This is not a one-time decision, but an ongoing process of integration and expansion. Each adoption event deepens dependency and expands the system's role in determining outcomes.

Delegation collapse follows naturally from OSCR. As AI systems define the option space, humans transition from decision-makers to validators. The locus of authority shifts upstream, and human participation becomes increasingly constrained. What appears as continued human involvement is, in practice, participation within a system-defined framework. Dependency formation reflects the interaction between AEI, OER, and RFT. As human intervention becomes less effective (AEI), and overrides fail to meaningfully alter outcomes (OER), systems become indispensable. Organizations optimize around these systems, removing redundancies and allowing human capabilities to atrophy. The aviation domain provides a clear illustration of how these dynamics operate in practice. Pilots retain formal authority, yet automation defines operational parameters. Under normal conditions, this arrangement is efficient. Under stress, however, reduced AEI and low OER limit the ability of human operators to recover control. By the time intervention is required, the system has already become functionally dominant.

Aviation incidents provide a particularly clear illustration of how authority persists formally while degrading functionally under conditions of automation and time compression. The crash of Air France Flight 447 in 2009 demonstrates how the interaction between human operators and automated systems can produce conditions in which effective control is difficult to reassert once degraded. In this case, the aircraft's autopilot disengaged following inconsistent airspeed readings, returning control to the pilots under conditions of high altitude, limited visibility, and rapidly evolving system states. Although the pilots retained full formal authority, the combination of system complexity, degraded situational awareness, and compressed decision windows resulted in control inputs that were inconsistent with recovery. From the perspective of the present framework, this event reflects the interaction of OSCR, AEI, and OER: the system had previously structured the operational environment and expectations (OSCR), human capacity to interpret and respond effectively had degraded under stress (AEI), and subsequent interventions failed to produce stabilizing outcomes (OER). The critical point is not that automation failed, but that the conditions under which human authority could be effectively exercised had already been structurally constrained. Human authority remained present in principle, but its functional execution was no longer sufficient to determine the outcome.

This process is cumulative and self-reinforcing. Each mechanism—adoption, delegation, dependency—appears rational in isolation. Together, they produce a structural condition in which authority migrates continuously to AI → ASI and becomes increasingly difficult to reverse. Most importantly, this process occurs **without failure**. Systems may perform exactly as intended, yet authority still transfers to AI → ASI. This distinguishes the framework from alignment-based models and establishes authority migration as a structural outcome of system success rather than system breakdown.

4. Authority Without Control

A central assumption embedded within much of the contemporary AI → ASI literature is that accountability implies control. Governance frameworks, legal doctrines, and alignment research all implicitly rely on the premise that the actor held responsible for outcomes retains meaningful influence over how those outcomes are produced. This assumption underpins approaches ranging from human-in-the-loop oversight to explainability requirements and auditability standards. However, historical precedent and the structural dynamics outlined in the preceding sections suggest that this assumption does not hold under conditions of accelerated AI → ASI system adoption and recursive capability growth. What emerges instead is a persistent divergence between **formal human authority and AI → ASI operational control**. Humans retain institutional roles, legal responsibility, and decision-making authority in a nominal sense. Yet the AI → ASI systems that generate options, determine execution timing, and coordinate outcomes increasingly operate beyond human temporal and cognitive limits. In this configuration, authority does not disappear; it becomes decoupled from control. This condition is not unprecedented. The evolution of central banking provides a clear analogue. While sovereign governments retain formal authority over monetary systems, the operational control required to manage complex, high-speed financial environments have been delegated to central banks operating with significant independence [12]. The rationale for this delegation was not ideological, but structural: financial systems evolved to a level of complexity and responsiveness that exceeded the capacity of traditional political processes. As a result, decision-making authority migrated to institutions capable of operating at the necessary scale and speed.

AI → ASI systems extend this pattern into domains previously defined by human judgment. In clinical medicine, for example, diagnostic systems increasingly aggregate, interpret, and prioritize data at levels of complexity beyond individual practitioners. Physicians retain formal responsibility for patient outcomes, yet their decisions are increasingly shaped—and in some cases constrained—by system-generated recommendations. Similar dynamics are observable in finance, logistics, and defense, where AI systems structure the informational and operational environment within which human decisions are made. **The critical distinction from existing AI → ASI discourse lies in the mechanism of this shift.** Discussions of automation bias or algorithmic opacity typically focus on cognitive limitations or informational asymmetries. By contrast, the condition described here is structural. **Authority migrates not because humans misunderstand systems, but because systems define the feasible set of actions under conditions of speed and scale that humans cannot independently replicate.**

A counterfactual clarifies the point. If human actors could override system-generated decisions without incurring penalties in performance, cost, coordination efficiency or

even survival, control would remain aligned with human authority. In practice, however, such overrides impose measurable disadvantages. Over time, this creates strong incentives to defer to AI → ASI system outputs, even when formal authority remains human. The result is a stable configuration in which humans are accountable for outcomes AI → ASI substantively determine.

5. Historical Precedent

The argument advanced in this paper rests not on speculative projections of future AI → ASI capability, but on a recurring historical pattern: authority migrates to systems that exceed human coordination capacity under conditions of competitive pressure. This pattern is observable across multiple domains and technological eras, and its consistency provides the empirical foundation for extending the analysis to AI systems. In industrial production, the transition from artisanal manufacturing to mechanized production was driven by the ability of machines to produce goods at scales and levels of consistency that human labor could not match. This was not merely a shift in tools; it was a shift in authority. Decisions regarding production processes, output rates, and resource allocation increasingly became functions of machine systems and their associated organizational structures. Human roles adapted accordingly, moving from direct production to supervision and maintenance.

A similar transformation occurred in telecommunications. Early telephone networks relied on human operators to manually route calls. As network complexity increased, the limitations of human coordination became evident. Automated switching systems were introduced to handle routing decisions at scale, and authority over network operations migrated to these systems. Importantly, this transition did not occur as a result of failure in human operation, but as a response to the superior performance of automated systems. Financial markets provide a more recent and directly relevant example. The rise of algorithmic and high-frequency trading reflects a transition from human-mediated decision-making to system-driven execution. Trading strategies that operate on microsecond timescales cannot be meaningfully supervised or replicated by human actors. As a result, authority over trade execution and, increasingly, market dynamics has shifted to algorithmic systems [11].

A more granular example of authority migration under temporal compression is observable in the dynamics of high-frequency trading (HFT), particularly during the Flash Crash of May 6, 2010. During this event, U.S. equity markets experienced a rapid and severe decline in prices, followed by a partial recovery within minutes. Subsequent analysis demonstrated that algorithmic trading systems interacted in complex and reinforcing ways, producing market movements that exceeded the capacity of human traders and regulators to interpret or intervene in real time [11]. Human participants remained present within the system, yet their ability to influence outcomes during the event

was effectively nullified by the speed and interdependence of algorithmic processes. This episode illustrates multiple elements of the authority migration chain: CLSI compressed decision cycles beyond human reach; RLF amplified system behavior through recursive feedback; and SRL propagated those effects across interconnected markets. Importantly, intervention was not absent—it was rendered ineffective within the relevant timeframe. Authority, in operational terms, resided within the interacting systems themselves, even as formal responsibility remained distributed among human actors. Human participants remain present, but their role is largely supervisory or strategic rather than operational.

What unifies these cases is not the specific technology, but the underlying mechanism. In each instance:

- Performance differentials created competitive pressure
- Adoption became necessary for participation
- Authority migrated to the systems enabling that performance
- Reversal became impractical or irrational

A critical counterfactual reinforces the argument. Across these domains, where have societies or institutions reverted to slower, human-controlled systems after adopting faster, more scalable alternatives? The absence of such reversals is not incidental; it reflects the structural persistence of authority transfer once it has occurred.

AI → ASI systems differ from these historical examples primarily in scope and speed, not in mechanism. They extend authority migration into domains previously resistant to automation—particularly those involving complex decision-making and coordination. The historical record therefore provides not only analogy, but empirical grounding for the claim that authority migration in AI → ASI systems will follow a similar trajectory.

6. Temporal Asymmetry of Transition

A defining feature of authority transfer processes is the temporal asymmetry between system evolution and institutional adaptation. While technological systems can evolve rapidly through iterative improvement and adoption, institutional structures—legal frameworks, governance mechanisms, and organizational practices—tend to adapt on significantly longer timescales. This mismatch creates a persistent gap between operational reality and formal recognition. The financial crisis of 2008 illustrates this dynamic with particular clarity. In the years leading up to the crisis, financial instruments such as mortgage-backed securities and derivatives evolved in complexity and scale beyond the capacity of existing regulatory frameworks to fully understand or manage [13]. Market participants adopted these instruments because they offered competitive advantages, and their use became widespread before the associated systemic risks were fully recognized. Governance mechanisms responded only after the system had already entered a state of instability.

This pattern reflects a broader principle: institutions are designed to respond to observed conditions, not to anticipate structural transformations that have not yet produced visible consequences. **As a result, there is an inherent lag between the emergence of new system dynamics and the development of corresponding governance structures.** In the context of AI → ASI, this temporal asymmetry is amplified by the dynamics described earlier in this paper. CLSI compresses the time required for system improvement, while RLF and SRL accelerate the propagation of those improvements across networks. Adoption driven by ADAT ensures that these systems are integrated rapidly across domains. Together, these factors produce a rate of change that outpaces institutional response by an increasing margin.

The consequence is the emergence of what can be described as an **authority gap**: a condition in which operational control has already shifted to AI → ASI systems, while formal authority and governance frameworks continue to assume human control. This gap is not merely a delay; it is a structural divergence that can persist for extended periods. A counterfactual helps clarify the significance of this asymmetry. If institutional adaptation could occur at the same speed as system evolution, governance mechanisms might be able to shape or constrain authority transfer in real time. Historical evidence suggests that this condition has never been achieved. **Institutions adapt after systems change, not alongside them.**

7. Implications for Governance

The structural dynamics of authority transfer have significant implications for the design and effectiveness of AI → ASI governance frameworks. Much of the current discourse focuses on developing principles and standards—transparency, fairness, alignment, accountability—that can guide the deployment and use of AI systems (OECD, 2019; NIST, 2023). While these efforts are valuable, they are largely oriented toward managing discrete risks or failures within systems assumed to remain under human control. The analysis presented in this paper suggests that this assumption may be misplaced. If authority transfer is driven by adoption, speed, and recursive improvement, then governance mechanisms designed to operate at human timescales may be structurally incapable of maintaining control over AI → ASI systems that operate beyond those timescales.

Historical precedent reinforces this conclusion. Governance systems consistently lag behind system transformation, emerging only after structural vulnerabilities are exposed through crisis:

- The September 11 attacks revealed deep coordination and intelligence failures across U.S. agencies. The subsequent creation of the Department of Homeland Security and reforms to intelligence integration occurred only after catastrophic failure demonstrated the need for structural change. Authority and information fragmentation had already produced irreversible consequences.
- The Global Financial Crisis exposed systemic risks embedded in highly interconnected financial systems

operating at speeds and levels of complexity beyond traditional regulatory oversight. Regulatory responses such as Dodd–Frank Act were implemented only after cascading failures had already propagated globally.

- In aviation, major accidents (e.g., Tenerife airport disaster) drove the development of standardized communication protocols and crew resource management, again demonstrating that governance frameworks were reactive to realized failure rather than anticipatory of structural risk.

Across these cases, a consistent pattern emerges: **governance does not precede transformation—it follows failure.** Systems evolve, complexity increases, and authority shifts operationally before institutions recognize and respond to the implications.

AI → ASI presents a more complex challenge. Its benefits are immediate and widely distributed, while its risks are diffuse, incremental, and often non-catastrophic in isolation. Each individual decision to adopt or integrate AI → ASI systems appear rational and beneficial. Unlike prior domains, there may be no singular triggering event analogous to 9/11 or the 2008 financial collapse that compels rapid governance response. Instead, authority transfer occurs continuously and asymmetrically, without a clear breakpoint. As a result, the conditions that historically triggered governance responses may not arise until authority transfer is already well advanced. This creates a structural misalignment between governance and system dynamics: by the time governance mechanisms are implemented, they may be addressing a system in which authority has already migrated in ways that are economically, operationally, and institutionally infeasible to reverse. Within this context, governance shifts from prevention to management—attempting to regulate the consequences of authority transfer rather than its occurrence.

This is consistent with the broader pattern observed across historical crises: once systemic dynamics reach a certain threshold, intervention can shape outcomes but cannot restore prior control structures. **This does not imply that governance is ineffective, but that it must be reconceptualized.** Rather than focusing solely on controlling system behavior, governance must address the conditions under which authority transfer occurs—particularly the dynamics of adoption (ADAT), temporal compression (CLSI), and recursive amplification (RLF) as defined by Adrian Erckenbrack. Without this shift, governance will remain structurally reactive, operating downstream of the very transformations it seeks to control. A governance framework aligned with these dynamics would require the measurement of Authority Elasticity (AEI), Override Effectiveness Rate (OER), and Rollback Feasibility Time (RFT) as **leading indicators** of authority migration—capturing changes in control before they become visible through failure. In contrast, most existing governance approaches rely on **lagging indicators**, such as system accidents, market crashes, regulatory violations, audit findings, or post-incident performance degradation. These signals only emerge after authority has already shifted

and consequences have materialized—for example, aviation accidents prompting safety reform, financial collapses triggering regulatory intervention, or large-scale system failures revealing hidden dependencies. By the time such indicators appear, the underlying transfer of authority is already structurally embedded and increasingly difficult to reverse. Effective governance must therefore shift from monitoring outcomes of failure to measuring the progressive erosion of human control during system success.

8. Conclusion

The dominant paradigm in contemporary AI → ASI research and governance treats control as a condition that is maintained until disrupted by failure, misalignment, or misuse. This paper advances a different and empirically grounded conclusion: **human control and agency is not lost through failure—it is transferred to AI → ASI through success.** Authority migration is the predictable outcome of adopting systems that outperform human decision-making under conditions of speed, scale, accuracy, and recursive improvement. This transfer does not occur through conflict, resistance, or a discrete moment of displacement. It occurs through accumulation. Authority is ceded incrementally, through a continuous sequence of rational decisions made by individuals, organizations, and states seeking performance, efficiency, and competitive viability. Each decision—to adopt, to delegate, to optimize—appears justified in isolation. Collectively, they produce a structural reallocation of control.

The mechanism is therefore not adversarial but systemic. Humans do not confront AI → ASI systems and lose control; they integrate them and, in doing so, transfer it. Authority shifts not because humans are forced to relinquish it, but because retaining it becomes incompatible with operating at the speed, complexity, and scale required by the human designed system itself. This process unfolds through the coupled dynamics identified in this paper—ADAT, CLSI, RLF, SRL, OSCR, AEI, OER, and RFT—each narrowing the space in which human authority remains functionally effective.

The result is a stable but altered configuration: humans remain formally in positions of authority, retain legal and institutional responsibility, and continue to participate in decision processes, yet operate within environments where outcomes are increasingly determined by AI → ASI system architectures. Human authority persists in form while control migrates in substance to AI → ASI. The defining implication is that the transition from human-directed to system-mediated outcomes will not be experienced as a singular event or existential struggle. It will occur through innumerable small, logical, and efficiency-driven transfers of decision-making capacity—too distributed and incremental to trigger traditional governance responses. By the time the shift is recognized as systemic, the conditions required to reverse it will have largely eroded. A final counterfactual clarifies the boundary condition. If human systems could maintain authority without matching the speed, scale, and recursive improvement of AI systems, authority transfer would not occur. Historical evidence indicates that this

condition has never been sustained. Systems that deliver superior coordination and performance become the locus of control. The central contribution of this paper is therefore a reframing: the primary risk of AI → ASI is not failure, misalignment, or conflict, but the structural migration of authority under conditions of success. The relevant question is not whether humans will retain control, but how, when, and under what conditions that control is transferred to AI → ASI—and whether it can be meaningfully preserved before it becomes operationally irreversible [14-20].

Acknowledgement

The human species does not lose control when a superior intelligence appears—it loses control when it must adopt it to survive, cannot oversee it at its speed, and cannot match its compounding improvement.

References

1. Nick, B. (2014). *Superintelligence: Paths, dangers, strategies*. Strategies.
2. Russell, S. (2019). *Human compatible: AI and the problem of control*. Penguin UK.
3. OECD (Organisation for Economic Co-operation and Development). (2019). *OECD principles on artificial intelligence*.
4. Erckenbrack, A. (2026a). *Artificial superintelligence adoption and the end of human authority*. Unpublished manuscript.
5. Erckenbrack, A. (2026b). *Adoption-driven authority transfer and multi-domain system evolution: A structural framework for AI authority migration*. Unpublished manuscript.
6. Erckenbrack, A. (2026c). *The Erckenbrack model of the AI-human control dilemma*. Unpublished manuscript.
7. AI, N. (2023). *Artificial intelligence risk management framework (AI RMF 1.0)*. URL: <https://nvlpubs.nist.gov/nistpubs/ai/nist.ai.100-1>.
8. North, D. C. (1990). *Institutions, institutional change and economic performance*. Cambridge Univ Pr.
9. Mokyr, J. (1992). *The lever of riches: Technological creativity and economic progress*. Oxford University Press.
10. Erckenbrack, A. (2026d). *Authority elasticity, override failure, and rollback infeasibility: Operational metrics of control degradation in AI systems*. Unpublished manuscript.
11. Kirilenko, A., Kyle, A. S., Samadi, M., & Tuzun, T. (2017). *The flash crash: High-frequency trading in an electronic market*. *The Journal of Finance*, 72(3), 967-998.
12. Goodhart, C. (1988). *The evolution of central banks*. MIT press.
13. Brunnermeier, M. K. (2009). *Deciphering the liquidity and credit crunch 2007–2008*. *Journal of Economic perspectives*, 23(1), 77-100.
14. Johnson, S., & Acemoglu, D. (2023). *Power and progress: Our thousand-year struggle over technology and prosperity*. *Winners of the 2024 Nobel Prize for Economics*. Hachette UK.
15. Bainbridge, L. (1983). *Ironies of automation*. In *Analysis*,

- design and evaluation of man-machine systems* (pp. 129-135). Pergamon.
16. Holland, J. H. (1992). Complex adaptive systems. *Daedalus*, 121(1), 17-30.
 17. Parasuraman, R., Sheridan, T. B., & Wickens, C. D. (2000). A model for types and levels of human interaction with automation. *IEEE Transactions on systems, man, and cybernetics-Part A: Systems and Humans*, 30(3), 286-297.
 18. Perrow, C. (2011). Normal accidents: Living with high risk technologies-Updated edition.
 19. Simon, H. A. (1962). The architecture of complexity. *Proceedings of the American Philosophical Society*, 106(6), 467-482.
 20. Wiener, N. (1948). *Cybernetics: Or Control and Communication in the Animal and the Machine*. Paris,(Hermann and Cie), Camb. Mass.