

FROM ARTIFICIAL INTELLIGENCE TO SUPER INTELLIGENCE

Intelligence is not *Authority*

Delegated cognition, institutional authority, and the political economy of intelligence



Intelligence can become super. Authority does not become automatic.

An argument in four parts

Abstract & Executive Position

I

The Shift

- 01 Why the Word Changed
 - 02 From Technology to Economic Resource
-

II

The Separation

- 03 Intelligence Is Not Authority
 - 04 SI Deepens Delegated Cognition
 - 05 The Harness Matters More — and Is Not Causum
 - 06 SI Makes Independent Adjudication More Defensible
-

III

The Field and the Firm

- 07 Research Is Turning to Institutions
 - 08 Causum Must Also Change
 - 09 SI and the Political Economy of Intelligence
-

IV

The Architecture

- 10 A Clearer Division of Institutional Labour
 - 11 From Human-in-the-Loop to Authority-on-the-Loop
 - 12 A Stronger Definition of Causum
 - 13 Strategic Consequence
-

Conclusion · Evidence Position · Selected Sources

ABSTRACT

The move in public and executive language from *Artificial Intelligence* to *Super Intelligence* is not proof that a new technical class of machine has arrived. It is evidence of something else: capable machine systems are being reframed from software artifacts into *sources of intelligence* — to be allocated, delegated, operationalized, and converted into economic action.

That shift strengthens the Causum thesis on one condition: that a single distinction stays foundational. Greater capability expands what can be reasoned, predicted, or executed. It does not establish who may bind an institution, interpret governing intent, accept risk, create obligation, or decide what others may rely upon.

Intelligence and institutional authority are not the same property.

Causum should therefore be indifferent to whether the capability behind a result is called AI, AGI, SI, ASI, a human expert, or a hybrid. Its concern is the matter requiring determination:

AUTHORITY What was delegated?	ASSERTIONS What was produced?	EVIDENCE What supports it?	KNOWLEDGE What governs it?	RELIANCE By what process may it be relied upon?
---	---	--------------------------------------	--------------------------------------	---

EXECUTIVE POSITION

Recent U.S. executive action adopted *Super Intelligence* as preferred terminology in federal non-statutory communications, while retaining — for now — the statutory definition of artificial intelligence. Science still uses *superintelligence* in a far stronger sense: future systems whose capability substantially exceeds humans' across important domains. Causum need not settle that dispute. SI makes a more durable proposition visible instead.

Greater intelligence does not automatically confer *greater institutional authority*.

That proposition cuts both ways:

STRENGTHENED

Causum as independent decision-clearing infrastructure.

WEAKENED

Any Causum that is a dashboard, a runtime-control layer, or a human review step that only matters while machines trail human experts.

01

Why the Word Changed

Artificial Intelligence names intelligence by its origin. *Super* Intelligence names it by comparison — a level of capability relative to some reference point. That rhetorical shift matters before anyone agrees whether today's systems meet the stronger, historical meaning of the word.

ARTIFICIAL INTELLIGENCE

Names its *origin*.

Asks: what can this technology do?



SUPER INTELLIGENCE

Names its *rank*.

Asks: who may direct it, who may rely on it, and who bears what it does?

Fig. 1 — The change in vocabulary changes the question.

The new framing forces a different question. Not only *what can this technology do?* but *how should an ever more powerful source of intelligence be used* — *who may direct it, who may rely on it, and who answers for its delegated action?*

That is exactly the territory of the **Political Economy of Intelligence**: not merely the price of machine cognition, but the production, allocation, delegation, control, consequences, and value of intelligence inside institutions and markets.

The SI transition does not create the Causum thesis. It makes it impossible to miss.

From Technology to Economic Resource

The early AI market was organized around inputs: models, GPUs, tokens, APIs, applications. The agentic market is organized around production: agents, harnesses, tools, machine work, operational outcomes. An SI economy goes one step further. It treats intelligence itself as a productive resource.



Fig. 2 — The value chain of intelligence. As capability rises, weight shifts to the right-hand side.

Once intelligence is economically accessible, it raises the questions every scarce productive resource raises:

- i *Who produces it?*
- ii *Who owns access to it?*
- iii *Who can afford it?*
- iv *Who may delegate authority to it?*
- v *Who captures its value?*
- vi *Who bears its losses?*
- vii *Who decides whether its conclusions deserve reliance?*

Tokens still matter; they measure machine-resource consumption. But no enterprise buys tokens for their own sake. It buys intelligence to get work done, decisions improved, risk managed, outcomes realized. The more capable the machine, the more the right-hand side of the chain — judgment, outcome, value — decides what it is worth.

II

PART TWO

The Separation

03

Intelligence Is Not Authority

There is a weak case for AI adjudication: machines make mistakes, so humans must check them. It gets weaker with every model release. If Causum depended on machine error, better intelligence would eventually erase its market.

Causum rejects that foundation.

The strong case is institutional, not statistical. **Intelligence and authority are different properties.** A system could out-analyse every human on a matter and still have no standing to bind an institution. It could identify geopolitical, cyber, actuarial, clinical, contractual, or financial risk with extraordinary accuracy — and still lack the authority to accept that risk, terminate a supplier, approve a claim, satisfy a regulator, create an obligation, or commit resources.

Capability	≠	Authority	Execution	≠	Legitimacy
Evidence	≠	Determination	Intelligence	≠	Institutional judgment

Fig. 3 — Four separations that survive any increase in machine capability.




Fig. 4 — Capability compounds continuously. Authority moves only when it is granted, expanded, or withdrawn by those entitled to do so.

These separations hold even if future systems exceed human capability across many domains. In fact, the stronger machine intelligence becomes, the less defensible it is to ground governance in human cognitive superiority. What remains is the durable question: *who or what is legitimately authorized to determine the matter?*

04

SI Deepens Delegated Cognition

Causum calls it *delegated cognition*: an institution supplies an objective, context, authority, tools, constraints, and an expected outcome, and machine intelligence decides parts of the path between intention and result.

Today's agents already approximate it. Frontier providers ship agents that execute code, manage files, use tools, operate applications, and run longer jobs with less supervision. As capability grows, specifying every intermediate step becomes economically irrational. The operating relationship changes:




Fig. 5 — More capable intelligence receives objectives, not instructions — and the distance between intent and reality grows.

SI accelerates this. The more capable the intelligence, the more institutions delegate objectives rather than instructions. As the distance between intent and execution widens, the **Intent–Reality Gap** becomes more consequential, not less.

05

The Harness Matters More — and Is Not Causum

SI does not diminish the enterprise AI harness. It validates it. Cloud and frontier providers are industrializing everything around the model: agent loops, tools, context, memory, state, identity, permissions, policy, security, observability, evaluation, integration. That layer turns capability into reliable machine work.

It also rules out a whole class of Causum positioning. Agent identity, orchestration, tool control, basic policy enforcement, telemetry, context management, monitoring, generic AI governance — hyperscalers and frontier labs have every structural incentive to own these. Causum should not try.




Fig. 6 — Two layers, two questions. The harness ends in evidence; Causum begins with it.

The functions are complementary. *The stronger the harness becomes, the better the evidence Causum has to work with.*

SI Makes Independent Adjudication More Defensible

The quiet risk of the AI era is concentration. One producer may control the model, the execution environment, the policy layer, the logs, the evaluation — and the declaration that the resulting behaviour was appropriate. The more consequential intelligence becomes, the harder it is to treat that self-assessment as equivalent to independent reliance.




Fig. 7 — Producer self-assessment versus independent reliance.

SI policy discourse already distinguishes internal controls from independent external evaluation and institutional oversight. That is not an endorsement of Causum, and should not be presented as one. It is evidence of a broader structural concern: separating production from independent evaluation.

Causum carries that question down to the matter itself. It does not ask whether a model is safe in the abstract. It asks whether *this* result, action, assertion, or decision can be relied upon in *this* institutional context. Independent model evaluation and independent decision adjudication are adjacent but distinct problems — and SI raises the stakes of both.

III

PART THREE

The Field and the Firm

07

Research Is Turning to Institutions

Frontier research increasingly treats the institutional environment as part of the design problem. Google DeepMind has argued that work on superintelligence should study coexistence among adaptive actors, preserve human agency, and treat institutions as *design primitives* — rather than assume a model can be governed in isolation. [\[Google DeepMind\]](#)

Anthropic's multi-agent research likewise anticipates institutions that become human–AI hybrids, or in some cases predominantly agent-based. [\[Anthropic\]](#)

YESTERDAY'S FRONTIER QUESTION

How capable is the model?



TODAY'S FRONTIER QUESTION

How do intelligent actors
coordinate inside institutions?

Fig. 8 — The research signal.

Neither lab proposes Causum, and this paper does not suggest otherwise. The signal is the direction of travel: the frontier question is moving from model capability toward institutional coordination among intelligent actors. That is precisely the environment in which a common substrate for assertions, evidence, governing knowledge, authority, and determination becomes economically interesting.

08

Causum Must Also Change

A thesis that is strengthened by SI must also be disciplined by it. Three commitments follow.

8.1

Causum cannot depend on human cognitive superiority

The adjudicator cannot be permanently defined as a human checking a machine. Legitimate authority over a matter may rest with any of these:



THE INVARIANT Not the human. *Legitimate authority over the matter.*

8.2

Causum should not become an SI safety platform

SI will draw enormous investment into containment, alignment, runtime policy, identity, sandboxing, monitoring, and model evaluation. Causum should *consume* those systems as evidence sources, not compete with them. The producer of intelligence produces intelligence and execution evidence. Causum constitutes the matter through which institutional reliance is determined.

8.3

Causum should not adopt SI terminology uncritically

Administrative use of *SI* already diverges materially from the stronger scientific meaning of superintelligence. Causum should preserve AI, AGI, SI, and ASI as *source-specific* terms, and use *machine intelligence* or *delegated intelligence* as its own architecture-level abstraction.



The architecture then stays independent of political branding, model generation, provider vocabulary, and disputed capability thresholds.

SI and the Political Economy of Intelligence

If intelligence is treated as productive infrastructure, the political-economy case sharpens. The question is no longer whether society has access to powerful models. It is whether access to intelligence — *and to the institutional machinery that converts intelligence into legitimate action* — is concentrated.




Fig. 9 — Nominal access to machine intelligence is not equal access to trusted decision capability.

Smaller organizations may reach the same frontier intelligence and still lack everything around it. That preserves the capability inequality identified in Causum's broader position. Causum does not redistribute intelligence itself. It lowers the fixed institutional threshold for participating in the economics of intelligence — by externalizing access to adjudicable decision capability.

*Access to intelligence should not require
owning the full institutional machinery that
turns intelligence into trusted decisions.*

IV

PART FOUR

The Architecture

10

A Clearer Division of Institutional Labour

The SI transition sharpens who does what. Six layers, each with a distinct job:

06	Reliance	The determination supports action, assurance, contractual reliance, payment, escalation, or further autonomy.
05	Authority	Legitimate human, machine, or institutional authority decides: Conform, Diverge, or Refer.
04	Clearing · Causum	Constitutes the matter, associates authority, determines admissible evidence, exposes unresolved claims, routes to determination.
03	Knowledge · KPG	Assertions, relationships, provenance, governing knowledge, evidence, and competing claims.
02	Harness	Operationalizes intelligence: context, tools, identity, policy, execution controls, memory, observability.
01	Intelligence	AI, SI, human expertise, and other cognitive systems produce analysis, propositions, predictions, and possible actions.

Fig. 10 — The six-layer architecture. Causum occupies the clearing layer — between what is known and who may decide.

Read from the bottom up, the layers form a single chain:



Fig. 11 — From intelligence to institutional reliance.

SI does not break this chain. It increases the amount — and the capability — of the intelligence entering it.

From Human-in-the-Loop to Authority-on-the-Loop

Human-in-the-loop governance assumes a person must approve machine activity because the machine is not capable enough. As capability improves, that assumption grows more expensive — and harder to justify.

Causum supports **authority-on-the-loop** instead. The test is not whether a human approved every action. It is whether legitimate authority was represented, whether its exercise was evidenced, and whether consequential divergence could be detected, challenged, referred, or revoked.



Fig. 12 — Two models of oversight, and the four movements of justified autonomy.

This is compatible with increasingly capable SI and with Causum's principle of *justified autonomy*. Under it, stronger intelligence can receive greater autonomy *without acquiring unlimited authority*.

Speed is made possible *by governance*.

A Stronger Definition of Causum

The AI-era definition of Causum was *infrastructure for making consequential machine activity adjudicable*. The SI era calls for a broader and more durable one:

DEFINITION

Causum is independent clearing infrastructure through which evidence-bearing matters involving human or machine intelligence are constituted, adjudicated, and converted into *legitimate institutional determinations*.

SURVIVES

- better models
- AGI
- SI
- multi-agent systems
- human-machine institutions
- machine adjudicators
- changing providers
- changing political terminology
- commoditized AI governance

Not AI.	Not SI.	Not the agent.	The common object is <i>the matter requiring determination.</i>
--------------------	--------------------	---------------------------	---

Strategic Consequence

SI strengthens Causum in four ways — and removes one weak foundation.





Fig. 13 — What SI strengthens, and what it takes away.

The durable proposition is different:

No degree of intelligence, by itself, establishes legitimate authority.



CONCLUSION

ARTIFICIAL INTELLIGENCE

framed the machine
as a *technology*.

SUPER INTELLIGENCE

frames it as a
capability.

POLITICAL ECONOMY

frames that capability
as *economic power*.

Causum addresses what comes next. As intelligence grows more capable and more abundant, the hard institutional problem is no longer how to produce cognition. It is how to establish when cognition may legitimately become action, obligation, reliance, or authority.

The enterprise harness

governs the execution of intelligence.

The KPG

structures what is asserted and known.

Causum

constitutes and clears the matter.

Legitimate authority

determines what may be relied upon.

The stronger machine intelligence becomes, the more these distinctions matter. As intelligence becomes abundant, the scarce institutional resource becomes *justified judgment*. Causum exists to make that judgment reachable, evidential, and adjudicable.

Intelligence can become super.
Authority does not become automatic.

Three kinds of claim, kept apart

Established fact

U.S. executive policy has adopted SI terminology while presently retaining the existing statutory definition of AI.

External evidence

Frontier providers are industrializing agent-harness and governance infrastructure; research increasingly examines autonomous agents, multi-agent institutions, human agency, coexistence, and independent evaluation.

Causum thesis

These developments create a potential market for independent decision-clearing infrastructure connecting intelligence, evidence, governing knowledge, authority, and institutional determination.

The third proposition remains a thesis to be demonstrated through deployment and market evidence. It is not presented as established industry consensus.

SELECTED SOURCES

- 1 [White House — Inaugurating the Era of Super Intelligence](#)
- 2 [American Presidency Project — White House Accord on Super Intelligence](#)
- 3 [Google DeepMind — Superintelligence, coexistence, institutions, and human agency](#)
- 4 [Anthropic — Trustworthy agents and multi-agent systems](#)
- 5 [Amazon Web Services — Amazon Bedrock AgentCore Harness](#)
- 6 [OpenAI — Frontier and agent safety materials](#)