

The Infrastructure of Power: Finance and Artificial Intelligence

Mirko Bradley

ORCID: 0009-0005-1080-9799

mirko.bradley@gmail.com

Abstract

The last four decades of global capitalism have been defined by a single structural logic, that of power exercised through the control of capital allocation. Financial infrastructure (investment banks, asset managers, sovereign wealth funds) has dominated this logic, determining which projects, technologies, and geographies receive resources and which do not. This paper argues that a structural transition is underway. Artificial intelligence, initially a tool deployed within financial infrastructure, is progressively redefining the horizon inside which that infrastructure reconfigures itself. Finance does not disappear, nor is it absorbed by AI; what changes is the field of possibility within which financial actors think, formulate problems, and identify solutions. Drawing on empirical evidence from the period 2022–2026, including the financialization of compute as an asset class, the operational penetration of large language models across high finance, the concentration of venture capital toward foundational model layers, the migration of financial risk categories into computational logic, and the replication of pre-2008 systemic risk geometry in cognitive infrastructure, the paper identifies the structural mechanism of this transition and its implications for the distribution of global power. The analysis is conducted from the standpoint of someone moving between both domains, a position that proves not merely biographical but epistemically necessary, since the transition is visible only from within the intersection it itself produces.

Keywords: Financial infrastructure, cognitive infrastructure, artificial intelligence, capital allocation, foundational models, systemic risk, power, surveillance capitalism, compute, venture capital concentration, fintech.

1. Allocation as a Form of Power

Power operates through force and consent. Yet there exists a quieter, more durable modality: the control of what is distributed to whom, according to which criteria, and with which systematic exclusions. This intuition cuts across intellectual traditions that

rarely intersect. Braudel saw it in the long-term material structures that shape the realm of the possible before historical actors even begin to choose (Braudel 1979). Foucault described it as the government of conduct, the production of subjects who act within fields of possibility defined elsewhere (Foucault 1978). Zuboff has rearticulated it for digital capitalism, showing how the control of human behaviour rests upon the control of the infrastructures that render it legible and predictable (Zuboff 2019). Three distinct perspectives, one common structure: the most effective form of power is that which remains invisible because it coincides with the very conditions of possibility for everything else.

Understanding how this logic assumed its contemporary form requires tracing it through older structures. Control over trade routes in the medieval Mediterranean already constituted a sophisticated form of allocative power long before anything resembling modern finance existed. The Italian maritime republics, above all Venice, did not produce the silk or spices that flowed through their ports; they controlled the conditions of their circulation. They decided who could trade, on what terms, and at what cost. Braudel demonstrated that this form of dominance is structurally more stable than direct control of production precisely because it remains indifferent to the specific content of what is allocated (Braudel 1979).

The Dutch East India Company, founded in 1602, marked the first global institutionalisation of this logic. It separated risk from control: the former was distributed across a broad base of shareholders, while the latter was concentrated in a small governing board. This architecture prefigured, by three centuries, the structure of today's asset managers. Power resided not in the cargoes of pepper or cotton, but in the capacity to decide which routes would be opened, which markets served, and which competitors excluded. The violence that accompanied this expansion was real, yet instrumental: its ultimate purpose was the mastery of the conditions of possibility for global trade rather than territorial conquest per se.

Nineteenth-century London brought this logic to its first fully financial expression. The gold standard, the market for sovereign bonds, and the clearing houses that regulated international payments transformed control over the flow of money into

a form of global governance exercised without armies or direct territorial control. A country unable to raise capital in London could not build railways, modernise its armed forces, or sustain its public debt in times of crisis. Control over the allocation of capital became synonymous with control over national development paths. This is the deep structure inherited, amplified, and globalised by twentieth- and twenty-first-century financial capitalism.

The 1980s did not invent this logic; they radicalised it in scale and velocity. Financial deregulation (initiated under Reagan in the United States and Thatcher in the United Kingdom, epitomised by London's 1986 Big Bang) dismantled the postwar constraints on capital mobility. The 1999 repeal of the Glass-Steagall Act in the US completed a process already underway in European markets. Capital was liberated from national borders and from the productive destinations that regulators had once tried to impose. It became fluid, global, and capable of moving in real time toward the highest expected returns, irrespective of geography or sector.

Two innovations in this period permanently altered the morphology of allocative power. First, the construction of financial indices as instruments of passive governance. When a fund replicates the S&P 500 or MSCI World, it does not exercise discretionary judgement; it delegates that judgement to the algorithm that defines the index. Those who design and maintain these indices thus wield enormous allocative power without formal authorisation or democratic legitimacy (Braun 2021). Second, the explosive growth of derivatives markets, which decoupled risk from underlying assets, enabled speculation on price movements without ownership, and generated exposures many times larger than the GDPs of entire nations. In redistributing risk, derivatives reshaped the architecture of control itself.

The resulting form of allocative power possesses a distinctive property that renders it exceptionally stable: it is distributed enough to be politically unassailable, yet concentrated enough to be structurally decisive. BlackRock does not manufacture cars, drugs, or energy, yet its investment committees influence the trajectory of entire industrial sectors worldwide. Gulf sovereign wealth funds do not govern foreign territories, yet their capital-allocation decisions redraw the geographies of global

development. Vanguard, State Street, and BlackRock together constitute the largest shareholder bloc in most major Western listed companies, an unprecedented concentration of ownership in the history of industrial capitalism (Braun 2021).

The causal chain linking this concentration to the daily lives of billions is so long that it becomes invisible. The pension contributions of millions of workers flow into the same vehicles that determine which firms employing those workers can access capital. When the Italian state issues debt to cover INPS deficits, those bonds are purchased on international markets where BlackRock is often a major buyer. The Italian pensioner, the American fund manager, and the asset manager's investment committee are connected by a real, material chain, yet none perceives the full circuit. This structural opacity is not the result of deliberate concealment; it is an emergent property of the system's architecture.

As Foucault showed, modern disciplinary power functions not primarily through repression but through normalisation: the construction of a field of possibility in which certain behaviours become natural and others unthinkable (Foucault 1975). Contemporary financial allocative power operates through an analogous logic. It does not explicitly forbid; it determines which projects receive funding, which technologies are developed and which remain speculative, which geographies are integrated into global circuits and which are left peripheral. The field of possibility is shaped before actors begin to choose, and it is shaped by those who control the flow of capital.

Zuboff extended this analysis to surveillance capitalism, demonstrating how digital platforms added a new layer: control not only over flows of money but over flows of attention and behaviour, turning human experience into raw material for prediction and modification (Zuboff 2019). This raises the central question of the present paper. If allocative power has historically migrated from trade routes to financial capital, and from financial capital to digital behaviour, what form is it assuming now? The argument advanced here is that artificial intelligence is not merely the latest instrument of existing allocative power. It is redefining the very conditions within which that power operates, producing a structural transition that no single actor has designed, but which emerges from the evolving geometry of the system itself.

The progressive dematerialisation of power is a historical constant, not a peculiarity of digital capitalism. Each shift toward greater abstraction generates the same illusion: that invisibility equals diffusion. Medieval trade routes were visible and therefore contestable. Modern financial credit was already more opaque. Digital platforms are more opaque still. Those who control the infrastructural layers do not merely channel information; they define the conditions under which information, choice, and action become possible. Artificial intelligence represents a further movement in this trajectory: it no longer merely shapes the channels through which information circulates, but the cognitive conditions within which information is generated, framed, and acted upon. The next section examines how this shift began.

2. Artificial Intelligence as an Instrument

The entry of computational technology into high finance represented an acceleration of the allocative power configuration described in the previous section. Financial infrastructures adopted technology in an instrumental manner, performing at greater speed, with higher precision, and on a broader scale what they were already doing. In this phase, which spans roughly the three decades from the mid-1980s to the mid-2010s, artificial intelligence functioned as a transparent instrument with respect to financial objectives. It amplified the operational capacity of the system without altering its underlying logic. Understanding this phase is essential, because the transition examined in this paper cannot be grasped without recognising its point of departure.

The broader frame for this acceleration is the world-system described by Wallerstein. Capitalism has historically organised itself through centre-periphery relations mediated by the control of commercial and financial flows. Centres extracted value from peripheries through mechanisms that changed form over time while preserving the same underlying geometry. When computational technology entered finance, it did so within this already consolidated structure. Computing and data-processing hubs emerged precisely in the places where capital had long been concentrated, Wall Street and the City of London. The periphery remained peripheral even in the algorithmic era. As Wallerstein anticipated for technological innovation

more generally, computational technology was absorbed by the centre of the world-system and used to reinforce its position rather than redistribute power (Wallerstein 2004).

The first major domain of penetration was algorithmic trading, which gained momentum in the 1990s and reached maturity in the following decade. Financial markets constantly generate price discrepancies between correlated instruments. Those who identify and exploit them first capture the profit. In the pre-algorithmic era this task fell to human traders who monitored markets and executed orders manually. Speed was the primary constraint. Algorithms removed that constraint. In the 2000s high-frequency trading pushed this logic to its extreme. Orders were executed in microseconds, positions were opened and closed thousands of times per day, and competitive advantage was measured in nanoseconds of latency. Citadel, DE Shaw and Virtu Financial built technological infrastructures whose operational complexity rivalled that of the largest investment banks. Physical proximity of servers to exchanges became a strategic variable as critical as the quality of the algorithms themselves (Patterson 2010).

Parallel to algorithmic trading, quantitative portfolio management emerged. Renaissance Technologies, founded by James Simons in 1982, remains the paradigmatic case. A fund staffed almost entirely by mathematicians, physicists and computer scientists, it employed statistical models to detect market patterns invisible to traditional analysis. Its flagship Medallion Fund generated annualised gross returns above 60 percent for more than three decades, a level of performance no discretionary strategy has ever matched. This outcome demonstrated that systematic inefficiencies exist in markets and can be exploited with sufficiently sophisticated quantitative tools, at least until widespread adoption erodes the edge. Emanuel Derman, a theoretical physicist who joined Goldman Sachs in the 1980s, documented from within this cultural shift. He showed how the language of theoretical physics progressively colonised the vocabulary of risk management (Derman 2004).

A third domain was risk management. Value at Risk (VaR) models, developed by JPMorgan in the 1990s and later adopted as an international regulatory standard

through the Basel accords, represented the first systematic attempt to quantify a portfolio's exposure to adverse market movements. These models estimate the maximum probable loss over a given time horizon at a given confidence level, based on historical return distributions. Their global adoption transformed risk management from a judgement-based practice into an algorithmic procedure. This standardisation produced an unanticipated effect. When all major institutions rely on similar risk models, their behaviours become correlated. Correlated behaviours amplify systemic volatility rather than reduce it, because all actors tend to sell the same assets in moments of stress.

A fourth domain was credit scoring and the automated evaluation of creditworthiness. From the 1980s onward, statistical models gradually replaced the discretionary judgement of bank officers in assessing loan applications. The FICO score, introduced in the United States in 1989, condensed an individual's credit history into a single number that determined access to credit and its terms. This automation standardised and accelerated evaluation, lowering operational costs and simultaneously crystallising the structural inequalities embedded in historical data, producing systematic exclusions that appeared neutral because they were mediated by algorithms. As MacKenzie has shown, financial models do not merely describe markets but actively perform them, transforming reality the moment they are adopted at scale (MacKenzie 2006). Credit scoring offers a clear example of this performativity.

The 2008 financial crisis marked the point where the narrative of the transparent instrument collided with the systemic reality it had helped create. VaR models had distributed risk through complex derivatives such as CDOs and CDSs, generating the illusion that risk had been eliminated through dispersion. In reality, it had been concentrated invisibly in the balance sheets of the most exposed institutions, masked by algorithmic ratings produced with the same models used to construct the instruments. When the US housing market deteriorated in 2007, the systemic correlation induced by uniform model adoption turned a sectoral crisis into a global collapse. The very mechanisms that had enabled expansion, the same logic of speed, correlation and risk distribution, were what amplified and accelerated the contagion.

Gorton documented how the crisis propagated through the shadow banking system, revealing that models had not dispersed risk but rendered it opaque (Gorton 2010).

The flash crash of 6 May 2010 added another dimension. Within minutes, US equity indices lost nearly ten percent of their value before recovering almost entirely in the same session. Investigations revealed that interactions among different trading algorithms, each reacting to signals generated by the others, had produced a self-reinforcing spiral of selling that no human operator had initiated or could stop. The system had developed an autonomous dynamic beyond the control of its designers. Michael Lewis described this period as the moment when Wall Street ceased to be a market where humans buy and sell assets and became instead a technological infrastructure in which algorithms interact with other algorithms at speeds inaccessible to human perception (Lewis 2014).

Across all these domains, the structure of power remained structurally unchanged from the configuration outlined earlier. Those who controlled the algorithms were the same actors who already controlled capital flows. Technology did not democratise access to financial markets in any structural sense. Instead, it further concentrated competitive advantage in the hands of those with the resources to develop and maintain the most sophisticated models. The instrument served the end, and the end continued to be defined by the pre-existing financial logic.

The 2008 crisis and the 2010 flash crash nevertheless revealed that this transparency was partial and unstable. Models generated systemic effects that their designers had neither foreseen nor could control. These events opened a fracture in the narrative of the neutral instrument. Through that fracture, the next phase began to emerge. The following section examines the moment when the instrument ceased to be transparent to its ends and started to participate in defining them.

3. The Transition Line

The previous sections outlined two distinct phases: one in which allocative power operated through control of financial capital, and one in which computational

technology entered that structure as an instrument transparent to its ends. This section examines the moment when the relationship between the two began to invert. It is a structural transition identifiable through a set of empirical indicators that, taken together, reveal a clear direction. Artificial intelligence ceases to function as a tool used by financial infrastructures and becomes the horizon within which those infrastructures reorganise themselves.

The first indicator is the transformation of compute into an autonomous asset class. Until the late 2010s, computational capacity represented an operational cost for financial institutions: servers, data centres and network infrastructure were necessary expenses rather than strategic investment targets. From 2020 onward, and with sharp acceleration after 2022, this category dissolved. Nvidia rose from a niche gaming company to the second-largest firm in the world by market capitalisation in just three years. Sovereign wealth funds and major asset managers began allocating capital directly into computational infrastructure rather than only into firms that used it. In September 2024 BlackRock, Microsoft, Global Infrastructure Partners and MGX announced the Global AI Infrastructure Investment Partnership, targeting between eighty and one hundred billion dollars to build data centres and the energy systems required to power them. Larry Fink described the initiative as a long-term investment opportunity measured in trillions of dollars, employing the same language historically reserved for railways and electrical grids. The signal is unambiguous: financial capital now recognises that control of cognitive infrastructure constitutes the next frontier of allocative power.

The second indicator is the full operational penetration of large language models into high finance. This adoption does not follow the typical S-curve of technological diffusion. In 2025, ninety-five percent of alternative fund managers reported using generative artificial intelligence tools in their work, up from eighty-six percent in 2023 (AIMA 2025). JPMorgan Chase deployed an internal suite of language models to more than two hundred thousand employees, with over half using them multiple times daily. Morgan Stanley achieved a ninety-eight percent adoption rate among advisory teams. Bridgewater launched in July 2024 a fund that relies on machine learning as the primary basis for investment decisions, integrating models from OpenAI, Anthropic and

Perplexity with proprietary systems. Balyasny developed an analytical platform capable of processing five million documents and answering portfolio manager queries in minutes, reducing analysis time from two days to thirty minutes. Several funds now report that cloud and GPU costs rival prime-broker financing costs. When cognitive infrastructure acquires the financial weight of credit infrastructure, the parallel is no longer metaphorical.¹

The third indicator is the extreme concentration of venture capital in the foundational layer. Venture funding in artificial intelligence does not follow the usual pattern of diffusion across numerous applications. Instead, it concentrates at the base: the foundational models on which everything else is built. In 2025 artificial intelligence captured nearly half of all global venture funding, with OpenAI and Anthropic alone absorbing fourteen percent of total global venture investment (Crunchbase 2026). In the first three months of 2026, funding for foundational artificial intelligence firms reached one hundred seventy-eight billion dollars, twice the total for 2025 and nearly six times the total for 2024. Investors in this layer are securing positions in the infrastructure upon which all other applications rest, following the same logic with which financial capital once invested in the rights to build railways before the railways themselves existed.

The fourth indicator is the migration of financial risk categories into a computational logic. The clearest signals come from supervisory institutions. The Financial Stability Board has highlighted systemic vulnerabilities linked to dependence on third-party providers, concentration among those providers, market correlations induced by uniform model adoption, cyber risks, and challenges related to data quality and model governance (FSB 2024). The US Treasury has expressed concerns about the concentration of advanced model development in a handful of firms, noting systemic risks comparable to those of a cyberattack capable of disabling the entire industry. Systemic risk, a category originally developed to describe interbank credit fragility, is

¹Data on Bridgewater and Balyasny are drawn from: A. Rennison and K. Martin, “Hedge funds race to adopt AI”, Financial Times, 14 October 2024.

now being applied to dependence on cognitive infrastructure. Finance is rewriting its vocabulary of risk within a computational logic shaped by artificial intelligence.

The fifth indicator is the replication of the pre-2008 geometry of systemic risk within cognitive infrastructure. The most precise structural parallel with the 2008 crisis concerns not the numbers but the geometry of dependence. Before 2008, risk appeared widely distributed through complex derivatives yet concentrated invisibly in a few critical nodes. Today three or four providers control the foundational models used globally. Should most financial institutions depend on models from the same vendor, any disruption could trigger cascading effects. Homogenisation of artificial intelligence-driven trading strategies could similarly amplify market volatility (FSB 2024). The earlier risk was primarily one of liquidity. The emerging risk is one of cognitive accessibility: if a foundational provider fails or restricts access, entire sectors lose the capacity to formulate problems before they lose the capacity to act on them. This represents a structurally new form of systemic risk for which existing regulatory frameworks were not designed.

Taken together, these five indicators reveal a coherent direction that none would show with equal clarity in isolation (FSB 2024; AIMA 2025; Crunchbase 2026). Compute becomes an asset class because capital recognises cognitive infrastructure as the next site of allocative control. Language models penetrate high finance because they redefine which problems can be formulated and which solutions remain accessible. Venture capital concentrates at the foundational layer because control of that layer shapes the field within which all others operate. Risk categories migrate into computational logic because dependence on cognitive infrastructure has become a genuine systemic variable. The geometry of concentration replicates the pre-2008 pattern because the extractive logic of capital tends structurally toward concentration regardless of the substrate on which it operates.

There is, however, a deeper dimension these indicators do not fully capture when viewed separately. Artificial intelligence functions simultaneously as the cognitive horizon within which finance reorganises itself and as the investment object that finance produces. Finance funds the cognitive infrastructure that then redefines

finance. This reflexive loop differs fundamentally from any previous relationship between capital and technology, in which the instrument and its financier remained distinct categories. Here the two co-produce each other, and this co-production accelerates the transition in ways that no single actor, including those who finance it, can fully control.

This configuration does not yet have a consolidated name in the literature. The term *cognitive techno-finance* is proposed here to designate the emerging structure in which financial capital and cognitive infrastructure are no longer separable: they co-produce, co-define and converge toward the same architecture of allocative control.²

This convergence can be expressed in compact formal terms. Let F denote the set of financial infrastructures and $T(t)$ computational technology at time t . In the phase described in Section 2, $T(t)$ operated as an instrumental function within F . The relation took the form:

$$F \supset T(t),$$

where F contained $T(t)$ as its instrument. The transition documented in this section marks the moment when this relation began to invert and the reflexive loop transformed the inversion into reciprocal co-determination. The mature relation becomes:

$$F(t) \leftrightarrow T(t), \text{ with } \partial F / \partial T > 0 \text{ and } \partial T / \partial F > 0.$$

The bidirectional arrow replaces containment. $F(t)$ and $T(t)$ co-determine each other with positive derivatives: the growth of each feeds the growth of the other. Artificial intelligence redefines the field within which finance operates, while finance fuels the expansion of artificial intelligence that further redefines that field. It is this circular, self-amplifying structure that the next section examines in its implications for the global distribution of power.

²The term “cognitive techno-finance” is proposed here as an original terminological contribution. Terms currently circulating in the literature, such as “AI-native finance” or “cognitive finance”, describe the penetration of artificial intelligence into finance without capturing the structural dimension of reciprocal co-production between financial capital and cognitive infrastructure.

4. The Reconfigured Horizon

The transition described in the previous section does not produce a simple structure of power. The field that emerges from the fusion of financial capital and cognitive infrastructure has no identifiable centre, no stable hierarchy and no single actor that governs it consciously. It possesses instead a distributed and conflictual geometry in which actors with different resources, objectives and positions compete for control over distinct layers of the same infrastructure. Understanding this geometry requires identifying the main poles of the system and the asymmetric relations that connect them.

The first pole consists of the major asset managers. BlackRock, Vanguard and State Street collectively manage more than twenty trillion dollars and hold significant stakes in almost all major publicly listed firms in Western markets. They act simultaneously as relevant shareholders in foundational cognitive providers, financiers of computational infrastructure through vehicles such as the Global AI Infrastructure Investment Partnership, and custodians of global pension savings. Their strength lies not in any particular set of holdings but in their structural position: any configuration of the system necessarily includes them as central actors. This position grants them considerable political immunity, since any regulatory intervention that seriously damaged their portfolios would simultaneously harm the retirement savings of tens of millions of people. The formal distribution of the capital they manage is their most effective protection.

The second pole comprises the foundational cognitive providers. OpenAI, Anthropic, Google DeepMind and Meta AI control the models on which everything else is built. Their advantage is technical yet structurally unstable: architectures evolve, benchmarks shift and dominance at the foundational layer today may disappear with the next technical transition. Their real power stems not from specific models but from control over training data, evaluation infrastructures and scientific talent networks that make external replication extremely difficult. They shape the cognitive horizon most directly: their technical choices, which tasks are optimised, which benchmarks are

adopted and how alignment is defined, determine which problems are formulable and which remain outside the field of the thinkable with currently available tools.

The third pole includes the technology giants as owners of the substrate infrastructure. Microsoft, Amazon and Google are not merely investors in cognitive providers; they own the terrain on which those providers operate. Cloud computing, data centres and global distribution networks are controlled by these three actors to such an extent that no foundational model can be trained or deployed at global scale without passing through their infrastructure. Their position resembles that of the casino owner: they need not be the most skilled players at the table, but the game is played on their premises. This position is structurally more stable than that of the cognitive providers because it remains indifferent to the specific content of the models running on their infrastructure.

The fourth pole consists of sovereign wealth funds and states acting as geopolitical players. Saudi Arabia, the United Arab Emirates, Singapore, Norway and China, through its state vehicles, bring enormous resources but pursue objectives that are not purely financial. They introduce logics of national sovereignty, geopolitical influence and the reduction of dependence on foreign infrastructures. The tension between the United States and China over advanced chips and foundational models represents the most visible expression of this dimension: control of cognitive infrastructure has become a national security variable before it is an economic one. These actors command resources others cannot match, yet they operate under political and institutional constraints that limit their agility.

The fifth pole includes major investment banks and hedge funds. JPMorgan, Goldman Sachs, Citadel and Renaissance Technologies possess the deepest experience of financial rules, yet they are discovering that those rules are shifting beneath them. They once held the broadest informational advantage. They now depend on cognitive infrastructures they do not control to maintain competitiveness. Their intensive adoption of large language models, documented in Section 3, is not merely a strategic choice but an adaptive response to the redefinition of the field in which they operate. The financial advisor working inside a system of pre-trained models has not lost their

role, but now exercises it within a horizon of problems already formulated, scenarios already selected and solutions already pre-filtered by the cognitive infrastructure. The space for genuine discretion narrows progressively without any deliberate decision to that effect.

The sixth pole comprises the regulators. The Federal Reserve, the ECB, the Financial Stability Board and the European Commission possess the formal authority to change the rules of the game, yet they remain structurally behind the evolution of the system they seek to regulate. The reports cited in Section 3 demonstrate that supervisory institutions have identified the systemic vulnerabilities arising from the concentration of cognitive infrastructure. However, existing regulatory frameworks were designed for liquidity and credit risks, not for risks of cognitive accessibility. The regulatory lag is therefore not only institutional but conceptual. Before effective intervention becomes possible, regulators must develop analytical categories adequate to this new form of systemic risk, a process that requires time the system does not grant.

What emerges from this map is not a pyramid with a clear apex but a multi-layered structure in which each pole exercises a distinct form of control over a distinct layer. Major asset managers govern the flow of capital. Cognitive providers shape the horizon of the thinkable. Technology giants control the physical substrate. Sovereign funds and states introduce the geopolitical vector. Banks and hedge funds remain the most experienced players yet also the most exposed to the redefinition of the field. Regulators can halt the game, but rarely before damage has occurred. None of these poles controls the system in its entirety. Each depends on the others for dimensions it cannot master autonomously.

In this configuration, the notion that power over the cognitive horizon is exercised by a conscious and identifiable subject proves to be a simplification. Negri and Hardt predicted in *Empire* (2000) that globalisation and digital networks would erode nation-state sovereignty and generate resistance from below, with the multitude emerging as a political subject capable of appropriating the very networks capital used to dominate it. What has occurred is structurally the opposite. Digital networks have concentrated power rather than distributed it, generated monopolies rather than

multiplicities, and the multitude has not appropriated the infrastructure but has become its raw material, as Zuboff demonstrated for surveillance capitalism (Zuboff 2019). Artificial intelligence carries this inversion to its most extreme point: the multitude does not merely supply the data on which models are trained. It operates within a cognitive horizon defined by those models, without any collective subject capable of contesting it, because contestation itself must occur with the tools that the horizon makes available.

The mechanism through which this power operates is not explicit prohibition but the construction of the field itself. Foucault showed that disciplinary power does not forbid but normalises, producing subjects who move spontaneously within boundaries they did not choose (Foucault 1975). The power of cognitive infrastructure functions through an analogous yet more radical logic: it does not produce disciplined subjects but generates the horizon within which any subjectivity takes shape. It is not Bentham's panopticon with a watcher from above. It is a structure that produces the horizon without any actor intending it, as the emergent outcome of thousands of technical choices made by researchers, engineers and product managers who never convened to decide what would become thinkable.

This paper forms the counterpart to the author's previous work, *The Trap* (Bradley 2026), which analysed the same phenomenon from the perspective of the population. That study modelled cognitive synchronisation as the progressive suppression of behavioural variance in populations exposed to algorithmic mediation infrastructures, predicting a threshold of supercriticality by 2050. The present paper examines the opposite side of the same transition: not the population that converges cognitively, but the infrastructure that produces this convergence and the poles of power that control it. The two works describe the two faces of one structural process: on one side the loss of cognitive variance, on the other the concentration of control over the horizon that reduces that variance.

The question this analysis leaves open, and which the concluding section will address, is not who exercises power over the global cognitive horizon. That question admits no single answer because the system has no centre. The relevant question is from

which position this transition can be observed in its entirety, and what that position implies for those who occupy it.

5. The Point of Observation

The question that closed the previous section deserves a direct answer. From which position can the transition described in this paper be observed in its entirety? The answer is not rhetorical. It is an epistemic argument with precise consequences for both the type of analysis this phenomenon requires and the credibility of those who attempt to conduct it.

Those who operate exclusively within financial infrastructures see technology as a tool because they inhabit the very techno-financial relation they analyse. From this vantage point, the logic of finance supplies the frame within which artificial intelligence appears as an accelerator, an optimizer, or a cost-reduction mechanism. The structural transition examined in this paper, the moment when the instrument begins to redefine the field, remains invisible from that position. Recognising it would require stepping outside the frame that the position itself provides. Those inside finance see AI through finance and, for that reason, cannot see the moment when finance begins to be seen through AI.

The symmetric position produces a parallel form of blindness. Those who operate exclusively within artificial intelligence research, model development or systems engineering tend to view financial infrastructures as an external context: merely the market that funds research without altering its internal logic. The field's deep dependence on capital allocation, the reshaping of research priorities by investment cycles, and the selection of benchmarks according to commercial rather than purely scientific criteria appear as given conditions rather than as variables internal to the system.

The position of the philosopher or social theorist who observes both fields from the outside generates a third form of blindness. Such observers can see structures but struggle to perceive them in motion. The very distance that enables systemic vision also

prevents the apprehension of speed and granularity. Established theoretical frameworks, from Foucault to Zuboff and from Braudel to Wallerstein, excel at describing consolidated structures but are less suited to capturing a structure while it is still forming. The transition analysed here is still underway. Its contours are therefore visible only to those close enough to perceive its movement yet distant enough to discern its direction.

There exists a fourth position, the one from which this paper has been written. It is neither that of the financier who has turned to technology, nor that of the technologist who has learned finance, nor that of the academic who studies both from afar. It is the position of someone who has deliberately constructed sustained mobility between the two domains, maintaining operational competence in both while preserving critical distance from each. This position is not a biographical accident but an epistemic condition. It determines what can and cannot be seen. The passage from the relation $F \supset T(t)$ to the relation $F(t) \leftrightarrow T(t)$ cannot be observed from one side of the system alone. It requires a standpoint that is simultaneously inside and outside, operational and reflective, technical and historical.

The professional figure that embodies this position does not yet have a consolidated institutional name, and this absence is itself significant. Roles receive names only when the system needs to replicate them at scale, that is, once the function has been institutionalised. The quant with systemic vision, the chief AI officer possessing genuine dual competence, the policy researcher with insider market experience, or the founder of cognitive infrastructure with a financial background represent partial approximations. The system has not yet learned to name this position because it has not yet finished forming it. Its competitive advantage lies precisely in the temporary window during which it remains rare.

It is important to distinguish this position from a superficially similar figure: the organic intellectual of power. The latter enjoys the most complete view of the system precisely because they occupy its highest levels. Keynes at the British Treasury during the First World War or Kissinger as architect of American foreign policy in the 1970s exemplify this type. They possessed exceptional systemic understanding and real

influence, yet their vision remained constrained by the logic of the system they served. They saw everything from within a pole and therefore could not perceive the structure of the poles themselves.

The epistemic position adopted in this paper is different. It consists not in the view from the centre of a pole but in mobility between poles. This mobility makes it possible to observe the asymmetric connections among them without being captured by the internal logic of any single one. It resembles the stance occupied by certain economic historians who moved fluidly between practice and theory: Fernand Braudel, who taught in Algeria before developing his historiography of the Mediterranean, or Albert Hirschman, who worked as a development economist in Colombia before theorising escape routes and unbalanced growth. Situated knowledge is not inferior knowledge. It is the necessary condition for perceiving phenomena that purely external or purely internal perspectives cannot reach.

This distinction carries a clear methodological implication. A paper written from inside a pole would tend to be prophetic: it would describe the system with insider precision and conclude with predictions of collapse or stabilisation. A paper written from the mobile transversal position is instead diagnostic. It describes a structure in motion, identifies the mechanisms driving it, and leaves open the questions the structure itself poses without claiming to resolve them from the outside. This is the kind of analysis that can be produced in 2026, when the transition is identifiable but not yet complete, when the contours of cognitive techno-finance are visible but its mature form is still emerging.

The point of observation is therefore not a marginal autobiographical detail. It forms an integral part of the argument. To claim that the transition from $F \supset T(t)$ to $F(t) \leftrightarrow T(t)$ is visible only from a mobile position between the two domains is to recognise that knowledge of this transition is structurally situated. It is not available from every vantage point but requires a specific position that carries a real cost of construction. This is not epistemic relativism. It is the acknowledgement that certain structural phenomena become visible only from certain positions, and that identifying the required position is part of the analysis itself.

The paper ends here, but the transition it describes continues. Cognitive technofinance remains a structure in formation. Its contours will continue to shift with the next generation of foundational models, with evolving geopolitical regulations, and with further concentration in the computational layer. What this paper has sought to capture is not the final form of that structure but the mechanism of its formation: the logic by which an instrument becomes a horizon, and the position from which that passage can be observed.

References

- AIMA (Alternative Investment Management Association) (2025). *AI Adoption in Alternative Asset Management*. London: AIMA.
- Bradley, M. (2026). *The Trap*. Zenodo. <https://doi.org/10.5281/zenodo.19609205>.
- Braun, B. (2021). “Asset Manager Capitalism as a Corporate Governance Regime.” In J. Hacker et al. (eds.), *The American Political Economy*. Cambridge: Cambridge University Press.
- Braudel, F. (1979). *Civilisation matérielle, économie et capitalisme, XVe–XVIIIe siècle*. Paris: Armand Colin.
- Crunchbase (2026). *Global Venture Report Q1 2026*. San Francisco: Crunchbase.
- Derman, E. (2004). *My Life as a Quant: Reflections on Physics and Finance*. New York: Wiley.
- Financial Stability Board (FSB) (2024). *Artificial Intelligence in Finance: Financial Stability Implications*. Basel: FSB.
- Foucault, M. (1975). *Surveiller et punir*. Paris: Gallimard.
- Foucault, M. (1978). *Sécurité, territoire, population*. Paris: Gallimard/Seuil.
- Gorton, G. (2010). *Slapped by the Invisible Hand: The Panic of 2007*. Oxford: Oxford University Press.
- Lewis, M. (2014). *Flash Boys: A Wall Street Revolt*. New York: Norton.

- MacKenzie, D. (2006). *An Engine, Not a Camera: How Financial Models Shape Markets*. Cambridge, MA: MIT Press.
- Negri, A. and Hardt, M. (2000). *Empire*. Cambridge, MA: Harvard University Press.
- Patterson, S. (2010). *The Quants: How a New Breed of Math Whizzes Conquered Wall Street and Nearly Destroyed It*. New York: Crown Business.
- Rennison, A. and Martin, K. (2024). "Hedge funds race to adopt AI." *Financial Times*, 14 October.
- US Department of the Treasury (2024). *Managing Artificial Intelligence-Specific Cybersecurity Risks in the Financial Services Sector*. Washington, DC: US Treasury.
- Wallerstein, I. (1974). *The Modern World-System I*. New York: Academic Press.
- Wallerstein, I. (2004). *World-Systems Analysis: An Introduction*. Durham: Duke University Press.
- Zuboff, S. (2019). *The Age of Surveillance Capitalism*. New York: PublicAffairs.