

The Epistemic Frontier of Capitalism: A Theory of Epistemic Regression through Artificial Intelligence and Misinformation

James Rice
University of Essex



© James Rice. This is an Open Access article distributed under the terms of the [Creative Commons Attribution Non-Commercial 4.0 License](https://creativecommons.org/licenses/by-nc/4.0/).

This article advances a single thesis about contemporary capitalism: under the pressure of frontier artificial intelligence and platform-mediated misinformation, accumulation, growth, and knowledge production has entered a phase in which outcomes are increasingly shaped by the privatization and strategic destabilization of our shared epistemic commons at the technological frontier. These, I define as consensus-based standards of scientific evidence, categories or criteria for a priori truth, and collaborative channels through which public reasons are disputed and agreed upon, these are facts and features of our politics that capitalism itself requires to coordinate expectations, channel investment, and establish and maintain political legitimacy. Artificial intelligence automates tasks and enhances productivity; but it also industrializes cognitive work itself, enabling the large-scale privatization of decision-making authority and the concentration of power, assets, and skills. AI defines our era, I suggest, but also threatens to disenfranchise a multitude of thinkers and creators, capable of contributing immense value.

Misinformation, on the other hand, threatens our sense of belonging in knowledge-based communities of practice, through the infusion of divisive, polarizing, and invasive falsehoods, intended to seed doubt, confusion, and anarchy. Misinformation, I argue, is the strategic manipulation of shared reality, often working at odds with reasonable intuition, and therefore deeply dangerous when undetected. The paper develops this thesis through two tightly linked case studies — AI and misinformation as strategies operating with the purpose of extreme capital accumulation funnelled to an elite subset of our technocratic society — and argues that capitalism now profits from corroding the epistemic and intellectual preconditions of its own vitality through information warfare.

Keywords: Epistemic Capitalism, Artificial Intelligence, Misinformation, Political Economy, Epistemic Commons, Economics of AI

Introduction: Capitalism's Epistemic Turn

Capitalism has always been, at its core, a system for coordinating uncertainty (Williamson, 1985; Shackle, 1972; Lachmann, 1986). From its earliest theoretical articulations, political economists understood capitalism not merely as a system of exchange or production, but as a social order that transforms dispersed, uncertain social knowledge into coordinated, collective, productive action. The price system aggregates dispersed information about scarcity and preference, firms serve to facilitate transactions when markets fail to coordinate efficiently, and states stabilize legal expectations by enforcing contracts, property rights, and macroeconomic order. After classical theory, subsequent literature in information economics and institutional theory further emphasized that uncertainty reduction is a precondition for efficient investment, innovation, and institutional legitimacy (Arrow,

1974; Hayek, 1945; North, 1990). What is distinctive about the present landscape is therefore not the existence of uncertainty as such, but the transformation of how uncertainty is governed, monetized, and politically contested under a new technological and institutional regime.

Two literatures have expanded rapidly over the past decade, but largely in parallel. One (our first case study) examines artificial intelligence, automation, and frontier technologies, focusing on productivity, labour markets, inequality, and long-run growth (Acemoglu and Johnson, 2023; Jones and Tonetti, 2020). The other (the second case study) analyses misinformation, fake news, and epistemic polarization, focusing on democratic stability, media ecosystems, and political behaviour (Vosoughi et al, 2018; Benkler et al, 2018). Despite addressing related phenomena, these literatures rarely speak to

one another in a systematic theoretical way. Artificial intelligence research often assumes a relatively homogeneous macroeconomic informational environment within which economic agents optimize, while misinformation research frequently treats technological change in information systems as a neutral transmission channel rather than as a reorganization of capitalist accumulation itself. In other words, technology mediates the flow of information, and misinformation corrupts or interrupts this flow. This paper will analyse these two phenomena in tandem, drawing on a third body of work, in analytic social epistemology and philosophy of technology, that, I argue, has recently developed the conceptual resources needed. These ideas and resources include models of hermeneutical injustice, accelerating technological change, and epistemic dispossession — ultimately, I integrate these concepts into a theory of technocratic capitalist accumulation driven, on two opposing (and sometimes allied) sides, by AI and misinformation (Milano and Prunkl, 2025; Alvarado, 2023; Kwok, 2025). This integration is ultimately what this paper contributes.

The separation between idealized epistemic ecosystems and unequal economic reality belies a deeper structural transformation. Artificial intelligence is not simply a labour-saving technology, and misinformation is not merely a tech-induced breakdown of civic norms. This paper argues that capitalism has entered a transformative phase in which the control of computational and cognitive inference—classification, prediction, ranking, and narrative framing—has become automated, concentrated, and central to accumulation. Under this regime, the production of shared knowledge is transformed from a public, professional function stabilized by institutional norms into an object of private investment, proprietary control, and mass strategic manipulation, embedded within platform-based business models and data-intensive physical infrastructures. Inference is becoming a core asset, and its privatization is reshaping accumulation, authority, and class relations at the same time.

The single thesis of this paper is that this transformation produces a structural contradiction specific to the present conjuncture: the more successfully capital encloses and destabilizes the epistemic commons, the more it erodes the conditions of its own long-run coordination. Epistemic frontier capitalism is, in this precise sense, self-undermining. It is a regime whose most profitable frontier is identical with the substrate of its own reproducibility. The argument proceeds by specifying the epistemic commons as a condition of coordination (Section 1), then developing the core claim that AI industrializes and encloses inference (Section 2), showing that misinformation is an endogenous accumulation strategy within this regime rather than an exogenous pathology (Section 3), and formalizing the resulting contradiction in dialogue with its classical antecedents in Marx and Polanyi (Section 4). Section 5 draws out the class and state implications, and the conclusion identifies the empirical signatures by which the regime can be recognized.

1. Capitalism, Coordination, and the Epistemic Commons

Any theory of capitalism at the epistemic frontier must begin with a definition of the economic structure itself. Capitalism is not reducible to markets, nor to private property alone. It is a historically specific social order characterized by generalized commodity production, private control over investment, competitive accumulation, and a state apparatus that enforces property rights while managing social conflict and macroeconomic instability. Crucially, capitalism also depends on institutions that induce stability, rendering the immediate future calculable and predictable, enabling economic and political actors to form expectations, coordinate plans, and assess the legitimacy of authority. Recent research in political economy emphasizes that the predictability of social structures is first and foremost social and based on dispersed and coordinated knowledge. Capitalism and democratic institutions rest on shared frameworks of valuation, expertise, and legitimacy that make coordination possible under radical uncertainty.

Marx's analysis foregrounded capital accumulation as the driving logic of capitalism, emphasising how surplus-value extraction structures class relations and is transformed by evolving technology. His Grundrisse formulation of the "general intellect" offers a sharper point of entry for the present argument than Capital does: it foregrounds how capital seeks to embed knowledge, coordination, and prediction into technological systems, altering the social distribution of agency (Marx, 1973). Where steam-era fixed capital embedded knowledge of mechanical processes, AI-era fixed capital embeds knowledge of social and cognitive processes. The general intellect now takes the form of trained weights, ranking functions, and recommender systems, and its enclosure is correspondingly more consequential because what is enclosed is not the labour of production but the labour of collective sense-making. Maximilian Kasy has recently argued that artificial intelligence is best understood as "the means of prediction," a productive asset whose ownership and governance determine the distribution of its benefits and whose privatisation therefore recomposes the underlying class relation (Kasy, 2025).

Weber approached capitalism from a complementary angle, emphasising rationalisation, bureaucracy, and calculability as fundamental cultural and organisational conditions (Weber, 1978). What Weber called calculability is one specific component of what I call the epistemic commons. Polanyi's contribution was to show that markets are not self-sustaining but depend on institutions that moderate commodification, and that the commodification of labour, land, and money generates systemic dislocation (Polanyi, 1944). Where Polanyi identified labour, land, and money as the "fictitious commodities" whose commodification threatens society, Jessop extends the idea of the epistemic commons as a fourth fictitious commodity (Jessop, 2007), one whose commodification threatens society in a deeper way because it is the condition of possibility of recognising the threats produced by the commodification of the other three.

What unites these canonical traditions is an implicit recognition that capitalism depends on

a socially agreed-upon epistemic order. Prices, contracts, and technologies only function when actors share sufficient agreement about facts, rules, and expectations; together with Keynes's observation that long-horizon investment rests on socially constructed expectations (Keynes, 1936; Beckert, 2016), regulatory standards, scientific expertise, and bureaucratic procedures historically transformed radical uncertainty into co-ordinable risk. The epistemic commons is the collective inferential substrate on which this scaffolding depends: the stock of shared vocabularies, categories, procedures of proof, and trusted channels through which agents interpret their circumstances and adjudicate competing claims.

Analytic social epistemology has recently developed the concepts needed to specify what happens when the basis for shared knowledge and discovery is undermined through adversarial shifts in the politically and technologically mediated information environment. Silvia Milano and Carina Prunkl, building on Miranda Fricker, suggest that there exists an ontological condition in which collective epistemic resources systematically prevent individuals from making sense of their social experiences (Milano and Prunkl, 2025). I argue that this very threat is being manifested through the imminent confluence of superintelligent AI and widespread online misinformation.

The object of this inquiry is the collective resources of intellectual labour, rather than information unique to the individual knower: a lacuna in the shared stock of interpretive resources is a deficiency in the biological medium within which agents form beliefs. Conversely, the epistemic commons is not a pure public good but a commons in Ostrom's sense (Ostrom, 1990), sustained by governance arrangements such as public debate over fundamental truths, professional editorial judgement, evidentiary rules of legal procedure, professional occupational licensing, financial and accounting conventions, statistical testing and disclosure standards, and the norms of scientific replicability and falsification (Popper, 1959). These arrangements are expensive to operate and vulnerable to displacement; that vulnerability is the opening through which the

accumulation dynamic identified in this paper moves.

The transition to the next section turns on a simple claim. These background conditions are not free, not automatic, and not self-reproducing. They are produced by specific institutions, under specific governance arrangements, and can be eroded by economic processes that do not internalize their value. The remainder of the paper analyses what happens when the frontier technologies of capitalist accumulation begin systematically to enclose and to degrade the very commons they presuppose.

2. AI as the Industrialization of Inference at the Epistemic Frontier

The first case study is artificial intelligence (AI). Technological change has always reshaped capitalism, but not all technologies are equal in their systemic effects. General-purpose technologies reorganise production, institutions, and social relations across the economy (Bresnahan and Trajtenberg, 1995). Artificial intelligence belongs to this category, but with a distinctive feature: it targets cognitive inference itself. Traditional industrial technologies mechanised physical labour and first-generation information technologies mechanised calculation and storage; artificial intelligence mechanises classification, prediction, and decision-making under uncertainty. In doing so, it transforms not only how goods are produced but how choices are made, authority is exercised, and epistemic legitimacy is constructed. Ramón Alvarado specifies the point: AI and the data-science methods associated with it are first and foremost epistemic technologies, designed to manipulate epistemic content through epistemic operations such as prediction and analysis (Alvarado, 2023).

Once we recognise AI as an epistemic technology whose functional role is to perform inference at scale, the central structural claim follows: the industrialisation of inference transforms a historically quasi-public function into a privately appropriable asset. Previous technological revolutions mechanised activities largely external to the commons; the AI revolution mechanises the intellectual activities

simultaneously internal to it — classification, adjudication, ranking, prediction — whose prior institutional carriers (courts, editorial boards, professional bodies, statistical agencies) were precisely the governance arrangements that kept the commons functioning.

AI systems increasingly substitute for human judgement in domains that involve interpretation, evaluation, and normative discretion. Coding algorithms, predictive financial systems, automated misinformation detection tools, and ubiquitous recommendation engines all perform functions once exercised by human agents embedded in institutional contexts with explicit accountability structures. By embedding these functions in algorithmic systems, firms and states relocate authority into technical, model-based artefacts that are difficult to contest or audit — what Pasquale diagnosed, well before the generative turn, as the “black box society” (Pasquale, 2015).

At the epistemic frontier, accumulation depends on owning and governing systems that define what content is visible, credible, and actionable: search engines, social media feeds, recommendation systems, and models that quantify risk. The economic value of these systems lies precisely in their capacity to shape expectations and behaviour at scale. Zuboff captures the velocity of the transition: “*surveillance capitalism vaulted to the commanding heights of the social order at astonishing speed*” (Zuboff, 2019), establishing a regime of what she calls “epistemic inequality.” Epistemic inequality is the distributional face of a deeper process in which the commons from which both parties to the inequality used to draw is itself being appropriated.

The enclosure has three interlocking consequences. First, it converts non-rival inferential capacity into an excludable asset: when classification and prediction are embedded in proprietary systems behind trade-secret protections, the evidentiary procedures that shape consequential decisions — creditworthiness, employability, newsfeed ranking, diagnostic triage — are withdrawn from public scrutiny. Chi Kwok frames this as “epistemic dispossession”: a structural form of

injustice in which socio-technical infrastructures systematically extract, suppress, or appropriate epistemic resources from those who supply them (Kwok, 2025). The user who contributes training data and the public whose discursive categories are reshaped by platform ranking are not exchanging epistemic goods with the firm on terms that leave the commons intact: they participate in an extraction whose principal output is a proprietary model and whose principal residue is a diminished shared resource.

Second, enclosure concentrates inference: data-intensive production exhibits increasing returns (Jones and Tonetti, 2020) and, combined with fixed costs of model training and scarce compute, produces structural concentration. A small number of firms own the dominant classificatory architectures through which large fractions of economic and political life are interpreted; concentration of inference produces something closer to ideological power than to ordinary rents. Third, and most consequentially, enclosure introduces a novel vulnerability. Proprietary inference systems generate profit by shaping the informational environment in which decisions are made; but the value of any particular inferential output depends on the user's belief that the underlying evidentiary process is reliable. When the same technology that generates credible outputs can also generate at scale what Fritts and Cabrera call "phantom patterns" — correlations that are statistically detectable but meaningless — the infrastructure of inference and the infrastructure of misinformation become physically and economically indistinguishable (Fritts and Cabrera, 2022). The commons is thereby not merely under-supplied by enclosure; it is actively polluted by the outputs that enclosure makes profitable to produce.

Productive inference expands the set of feasible goods and services and contributes to long-run growth. Extractive inference redistributes surplus by manipulating informational asymmetries, attention, and predictable consumer behaviours for the purpose of profit. Both generate surplus, but only the former aligns inference-accumulation with social welfare. In the regime described here, the systematic tendency is toward the latter

because it is the latter that most efficiently exploits the commons. The intellectual infrastructure of artificial intelligence — the mathematics and the science — once carried by public or quasi-public research-based institutions, has been privatised, concentrated, and rendered productively indistinguishable from its pathologies. We now move from AI to misinformation because agglomeration and the erosion of cultural integrity is the structural condition from which misinformation emerges as an accumulation strategy unique to platform economies mediated by advanced and ubiquitous forms of machine learning.

3. Misinformation as an Intentional Strategy for Asymmetric Accumulation

The second case study is misinformation. Misinformation is typically framed as a pathology of democratic discourse, a byproduct of cognitive bias or technological disruption. Within the framework developed here, it must be theorised as an endogenous outcome of accumulation rather than an external shock. When profits depend on controlling access to information, synthesising artificial expectations, and skewing belief formation, the strategic production of misleading information becomes economically rational, even as it imposes substantial social costs. Misinformation is not a glitch; it is an equilibrium feature of the system. Stiglitz and Kosenko have recently shown that standard results on market efficiency break down once disinformation is treated not as noise but as a strategic choice variable deployed by informed actors against less informed counterparties (Stiglitz and Kosenko, 2024a, 2024b).

False or misleading content diffuses differently from reliable information, often spreading more rapidly and widely within networked environments (Vosoughi et al, 2018). This pattern is not accidental. Platform architectures reward engagement, novelty, and emotional salience, creating incentives that favour content capable of capturing attention regardless of its veracity (Allcott and Gentzkow, 2017). Three mechanisms sustain the equilibrium. The first is the engagement economy itself. Platforms monetise attention, and attention responds disproportionately to affect, novelty, and

violation of expectation. It follows — from the business model rather than from any individual bad actor — that content optimised for engagement will tend to be content whose relationship to truth is incidental. The saturation is the emergent property of competition for attention on a substrate that does not reward the costlier signals that correlate with truth.

The second mechanism is the generative supply shock. Large language models and diffusion models have lowered the marginal cost of producing plausible-seeming content to approximately zero. The political-economic implication is that the ratio of synthetic-to-curated content in circulation will rise, not because anyone has decided it should, but because a firm that declines to use generative tools is outcompeted by one that does (Kay, Kasirzadeh and Mohamed, 2024). The structural shift from a curated-content equilibrium to a generative-content equilibrium changes what the epistemic commons contains. Under the curated equilibrium, even highly partial sources had to be produced by identifiable actors who incurred real costs and could be reputationally sanctioned. Under the generative equilibrium, vast quantities of content can be produced with no identifiable actor and no reputational exposure. The distinction between content and noise begins to erode because the cheapest path to communicative bandwidth runs through generation rather than through epistemic effort.

The third mechanism is the manufactured-doubt strategy, which predates the present moment but is amplified by it. Oreskes and Conway documented how tobacco and fossil-fuel interests discovered that it was cheaper to manufacture uncertainty about scientific consensus than to contest it on the merits (Oreskes and Conway, 2010). What is different now is that the infrastructure for manufacturing doubt is continuously available, inexpensive, and targeted, and asymmetric media ecosystems institutionalise the dynamic (Benkler et al, 2018; Jungherr, 2023). The manufactured-doubt strategy works by undermining the conditions of political self-determination at the commons level rather than at the claim level. A regulator persuaded that the

scientific consensus on tobacco or climate is contested does not need to be persuaded that any particular claim is false; the regulator has already been moved out of the range in which the question of truth is actionable. Scaling this effect through AI-mediated infrastructures is, from the strategic actor's perspective, not a corruption of the infrastructure but a use of it.

Beliefs function as coordination devices in capitalist societies. Shared expectations about monetary and fiscal policy, technological progress, the future of regulation, and social norms enable long-horizon investment and collective action (North, 1990). When the structure of public beliefs becomes fragmented, coordination costs rise. For actors seeking to delay regulation, avoid liability, or weaken collective action — for example on climate change — the fragmentation of public knowledge can be advantageous. Misinformation increases uncertainty and disagreement, reducing the probability that coherent political or regulatory responses will emerge. Economic actors do not need to believe misinformation themselves for it to be strategically useful; what matters is that misinformation disrupts common knowledge, so that citizens, regulators, and investors cannot agree on basic facts and policy regresses into a contested, delayed, ineffective state.

What unites these three mechanisms is that each generates private returns while depleting a shared resource. The firm that optimizes for engagement, the firm that floods a market with synthetic content, and the actor that manufactures doubt about climate science all capture value by increasing the costs of collective sense-making. Milano and Prunkl's concept of "epistemic fragmentation" names the cumulative result. The public does not merely receive different information under algorithmic mediation: they lose shared access to the interpretive context within which information can be evaluated (Milano and Prunkl, 2025). Fragmentation is not simply heterogeneity of belief. A heterogeneous public can deliberate when it shares evidentiary standards and a common vocabulary of adjudication; a fragmented public cannot, because the interpretive substrate needed to translate

between positions has been partitioned into incompatible proprietary contexts.

The basic argument and conclusion for this second case study is not that misinformation is qualitatively new or that audiences are uniquely credulous. It is the narrower structural claim that the cost function of epistemic commons-degradation has collapsed while the private returns to such degradation have risen — sufficient, in the elementary economics of externalities, to guarantee over-production without any pathology of individual actors. The over-production is not a side-effect of the accumulation process but the accumulation process itself: misinformation is what the system produces in the course of producing its intended output.

4. The Self-Undermining Contradiction

With the case studies of artificial intelligence and misinformation — Sections 3 and 4 — in hand, the contradiction can now be stated in its sharpest form. Capitalism requires a functioning epistemic commons in order to coordinate expectations, pricing, contracting, regulation, and political legitimacy. The frontier technologies through which contemporary capitalist accumulation occurs — AI-mediated inference and platform-mediated communication — generate returns that are maximized by enclosing and degrading those commons. The social accumulation of highly valued creative capital — the means of cultural production — in the hands of powerful capitalists actually erodes the structural conditions of its own reproduction, through the exclusion from power of ideological diversity based on bespoke knowledge gained through quotidian personal experience. This is the self-undermining contradiction that names the regime. It is first political and cultural, but it is also structural. It arises from the alignment of profit incentives with informational volatility, an alignment that becomes more pronounced as shifts in ideology and the personalization of the informational environment are together commodified.

This contradiction is distinct from the classical contradictions it resembles. Marx's tendency of the rate of profit to fall concerns the internal dynamics of surplus-value extraction. Polanyi's

double movement concerns the social destruction wrought by commodifying land, labour, and money. The substrate here — collective inferential capacity — is categorically different: it is not a prior material or social resource that markets transform into a commodity, but the condition of possibility of market coordination itself. Degrading labour produces immiseration; degrading land produces ecological crisis; degrading money produces financial panic. Degrading the epistemic commons produces a dissolution of the shared interpretive and epistemic ground on which any of these crises can be recognised, contested, and responded to at all.

Kohei Saito has argued that capital is structurally indifferent to the natural conditions of its own existence and is therefore unable to internalise the value of those conditions ([Saito, 2023](#)). The same applies to the epistemic case. Capital is indifferent to the epistemic properties of the outputs it produces, so long as those outputs generate engagement, adoption, or throughput. The difference between an inferential output that augments the commons and one that degrades it does not register at the level of the valorisation process; it registers only downstream, in the degradation of the social and relational character of capitalism itself.

Habermas's account of democratic legitimacy offers the complementary normative diagnosis. Democratic orders require a public sphere capable of generating communicatively validated consensus, and such a sphere presupposes discursive conditions that cannot be sustained when communication is systematically colonised by strategic action ([Habermas, 1996](#)). A political order whose communicative substrate has been colonised by strategic action does not merely produce worse deliberation; it loses access to deliberation as a political technology, and with it the means by which the political system ordinarily learns what it has gotten wrong.

Democratic systems presuppose a minimally shared environment for building trust — agreement on facts, procedures, and standards of evidence — within which disagreement can be productively managed. When belief formation is privatised and optimised for

engagement rather than truth, democratic deliberation becomes structurally vulnerable to manipulation. Political conflict shifts from distributive struggles over material outcomes to epistemic struggles over the nature of reality itself, making compromise increasingly difficult. The economic consequences are equally significant: instability raises uncertainty and weakens the informational signals that guide investment, so that short-term gains for some firms come at the cost of long-run growth and innovation.

Two objections deserve a brief response. The first objection is that market incentives themselves will resolve the contradiction between fundamental social and cultural value creation, on the one hand, and intangible capital accumulation without meaningful economic purpose on the other, as firms whose outputs are disconnected from physical assets or products lose credibility and revenue. This misunderstands the structure of the externality: an individual platform's engagement-optimised feed imposes epistemic costs diffusely while capturing attention-based revenue privately, with no mechanism to reprice the diffuse cost. The structure is isomorphic to classical commons dilemmas, in which private optimisation produces collective under-provision (Olson, 1965; North, 1990). Moreover, reputational sanction depends on the prior existence of a shared evidentiary standard against which claims can be adjudicated. Under fragmentation, that sanction becomes structurally unavailable.

The framework yields two substantive predictions. First, regulatory interventions focused on the content of misinformation will be less effective than interventions focused on the underlying infrastructure of digital media and commoditized intelligence — the governance of training data, model access, ranking systems, and compute — which serves to enable the pervasive dissemination of misinformation online. If, across jurisdictions, content-level interventions systematically succeed where infrastructure-level interventions fail, the thesis is wrong. Second, conventional measures of successful capitalist societies — accelerating productivity growth, strong profit margins, increasingly valuable market capitalisation —

may rise even as the systematic capacity for digital coordination and consensus declines. If these two measures move together, the thesis is wrong; if they diverge in the predicted direction, my argument, as it stands, would seem to be supported.

5. Class, State, and the Reconfiguration of Power at the Epistemic Frontier

Two further implications follow from the thesis. The first concerns class. Within the regime described here, the decisive means of production is inferential capacity, and the decisive class relation is the one structured by access to and control over that capacity. Those who own the (increasingly blended) physical and intellectual infrastructure of inference — factories for chips and compute, repositories of training data and model weights, experts who understand ranking systems, and CEOs who run and build hyperscaling platforms — exercise a structural power not fully captured by either ownership of physical capital or command of ordinary expertise. On the other side stand epistemic subjects: those whose cognitive and affective activity supplies the training data, whose behaviour is targeted for inferential intervention, and whose interpretive environments are constituted by outputs over which they exercise no governance. Inequality of inference compounds inequality of income and wealth, because the inferential tools most effective at preserving and growing wealth are themselves the most expensively governed and the most tightly held.

The second implication concerns the state. The state is now also the principal institution with the scale and authority to sustain the epistemic commons against enclosure. A state that is unable to effectively and transparently audit algorithmic systems, counter strategic misinformation, articulate credible public narratives, or produce authoritative statistical knowledge under conditions of synthetic-content saturation is not merely a weaker state. It is a state whose other functions — fiscal, regulatory, adjudicative — are progressively hollowed out, because all of those functions depend on the shared evidentiary substrate that the epistemic commons supplies. The comparative political economy of how different institutional variants respond is an important

agenda, but the structural diagnosis is prior to it: across variants, the counter-movement must do its work under conditions in which the shared evidentiary ground it needs is itself the object under attack.

The thesis does not entail that AI is, in any essentialist sense, opposed to the epistemic commons. Under different institutional arrangements — public training data, open model access, regulated ranking, compute governed as a quasi-utility — the same inferential technologies could be deployed in ways that augment rather than enclose the commons. The claim is structural, not technological: it concerns the accumulation dynamic produced when these technologies are governed by private, engagement-driven incentives under conditions of extreme concentration. Changing the technology without changing the governance will not resolve the contradiction; changing the governance without abandoning the technology can.

Conclusion: The Empirical Signature of Epistemic Frontier Capitalism

Capitalism's historical resilience has rested on its capacity to harness technological innovation while preserving sufficient order to coordinate investment, promote global governance, and maintain institutional legitimacy. From the rise of professional expertise to the institutionalization of scientific authority and bureaucratic rule, the development of capitalism has depended on shared systems of knowledge that transform uncertainty into valued expectations for economic continuity. Epistemic frontier capitalism represents a new phase in which this balance becomes increasingly fragile. AI industrializes the production of inference, converting prediction and judgment into scalable, proprietary assets. Platform-based economies monetize algorithms that define how individuals form beliefs, embedding profit incentives directly into the architecture of our digital information environments. Within this configuration, misinformation emerges not as an aberration but as a socially corrosive accumulation strategy in a system that rewards epistemic disruption.

The central contradiction is that accumulation increasingly depends on destabilising the informational foundations upon which capitalism itself relies. Markets, states, and firms require relatively stable expectations to coordinate long-term investment and collective action, yet the same technologies that enhance predictive power also enable the strategic manipulation of belief, the fragmentation of common knowledge, and the erosion of institutional credibility. The contradiction is a structural feature of the regime and will intensify as inferential technology becomes more capable, more concentrated, and more deeply integrated into public life. Whether epistemic frontier capitalism evolves toward renewed stability or deeper fragmentation will depend on whether inferential capacity can be governed as a quasi-public good, or whether epistemic power remains concentrated and democratic institutions struggle to adapt.

Classical and contemporary theories alike have often treated epistemic stability as a background condition rather than a contested terrain. The accumulation of epistemic power and the degradation of the epistemic commons are two faces of the same process. The question is no longer how to correct a market failure in the provision of "good information." The question is whether the political-economic order can sustain the conditions of its own coordination when the most profitable frontier of capital is the privatisation and destabilisation of collective sense-making. Answering it is, as Polanyi saw in a different crisis, ultimately a matter of whether society can re-embed a market that has begun to consume the substrate that makes society possible. Epistemic frontier capitalism poses a question that is at once economic, political, and philosophical: can a system that profits from epistemic disruption sustain the social order it requires to function? I have argued that unless the epistemic commons is reconstituted as an ecosystem for the development of collective and creative human livelihoods, unfettered technological capitalism will continue to accelerate in consuming the very ground on which our civilizations' moral and cultural foundations lie, for now, unprotected. It is our duty to continually seek out and defend these ultimate human values.

References

- Acemoglu, D. and Johnson, S. (2023). *Power and Progress: Our Thousand-Year Struggle Over Technology and Prosperity*. New York: PublicAffairs.
- Allcott, H. and Gentzkow, M. (2017). Social Media and Fake News in the 2016 Election. *Journal of Economic Perspectives*, 31(2), pp. 211–236.
- Alvarado, R. (2023). AI as an Epistemic Technology. *Science and Engineering Ethics*, 29(5), article 32. <https://doi.org/10.1007/s11948-023-00451-3>.
- Arrow, K.J. (1974). *The Limits of Organization*. New York: W. W. Norton.
- Beckert, J. (2016). *Imagined Futures: Fictional Expectations and Capitalist Dynamics*. Cambridge, MA: Harvard University Press.
- Benkler, Y., Faris, R. and Roberts, H. (2018). *Network Propaganda: Manipulation, Disinformation, and Radicalization in American Politics*. Oxford: Oxford University Press.
- Bresnahan, T.F. and Trajtenberg, M. (1995). General Purpose Technologies: Engines of Growth? *Journal of Econometrics*, 65(1), pp. 83–108.
- Fritts, M. and Cabrera, F. (2022). Online Misinformation and ‘Phantom Patterns’: Epistemic Exploitation in the Era of Big Data. *Southern Journal of Philosophy*, 60(1), pp. 57–87. <https://doi.org/10.1111/sjp.12451>.
- Habermas, J. (1996). *Between Facts and Norms: Contributions to a Discourse Theory of Law and Democracy*. Translated by W. Rehg. Cambridge, MA: MIT Press.
- Hayek, F.A. (1945). The Use of Knowledge in Society. *American Economic Review*, 35(4), pp. 519–530.
- Jessop, Bob. 2007. “Knowledge as a Fictitious Commodity: Insights and Limits of a Polanyian Perspective.” Reading Karl Polanyi for the Twenty-First Century: Market Economy as a Political Project, edited by Ayşe Buğra and Kaan Ağartan, 115–133. Basingstoke: Palgrave Macmillan.
- Jones, C.I. and Tonetti, C. (2020). Nonrivalry and the Economics of Data. *American Economic Review*, 110(9), pp. 2819–2858.
- Jungherr, A. (2023). Artificial Intelligence and Democracy: a Conceptual Framework. *Social Media + Society*, 9(3), pp. 1–14. <https://doi.org/10.1177/20563051231186353>.
- Kasy, M. (2025). *The Means of Prediction: How AI Really Works (and Who Benefits)*. Chicago: University of Chicago Press.
- Kay, J., Kasirzadeh, A. and Mohamed, S. (2024). Epistemic Injustice in Generative AI. Proceedings of the 2024 AAAI/ACM Conference on AI, Ethics, and Society. New York: ACM, pp. 684–697. <https://doi.org/10.5555/3716662.3716722>.
- Keynes, J.M. (1936). *The General Theory of Employment, Interest and Money*. London: Macmillan.
- Kwok, C. (2025). Epistemic Dispossession in Platform Capitalism: Insights from Iris Marion Young’s Five Faces of Oppression. *Social Epistemology*, published online 3 December 2025. <https://doi.org/10.1080/02691728.2025.2585445>.
- Lachmann, L.M. (1986). *The Market as an Economic Process*. Oxford: Basil Blackwell.
- Marx, K. (1973). *Grundrisse: Foundations of the Critique of Political Economy* (M. Nicolaus, Trans.). Penguin Books.
- Milano, S. and Prunkl, C. (2025). Algorithmic Profiling as a Source of Hermeneutical Injustice. *Philosophical Studies*, 182(1), pp. 185–203. <https://doi.org/10.1007/s11098-023-02095-2>.
- North, D.C. (1990). *Institutions, Institutional Change and Economic Performance*. Cambridge: Cambridge University Press.

Olson, M. (1965). *The Logic of Collective Action*. Cambridge, MA: Harvard University Press.

Oreskes, N. and Conway, E.M. (2010). *Merchants of Doubt*. New York: Bloomsbury Press.

Ostrom, E. (1990). *Governing the Commons: The Evolution of Institutions for Collective Action*. Cambridge: Cambridge University Press.

Pasquale, F. (2015). *The Black Box Society: The Secret Algorithms That Control Money and Information*. Cambridge, MA: Harvard University Press.

Polanyi, K. ([1944] 2001). *The Great Transformation: The Political and Economic Origins of Our Time*. Boston: Beacon Press.

Popper, K. R. (1959). *The Logic of Scientific Discovery*. Hutchinson.

Saito, K. (2023). *Marx in the Anthropocene: Towards the Idea of Degrowth Communism*. Cambridge: Cambridge University Press.

Shackle, G.L.S. (1972). *Epistemics and Economics: A Critique of Economic Doctrines*. Cambridge: Cambridge University Press.

Stiglitz, J.E. and Kosenko, A. (2024a). *The Economics of Information in a World of Disinformation: A Survey. Part 1: Indirect Communication*. NBER Working Paper No. 32049. Cambridge, MA: National Bureau of Economic Research.

Stiglitz, J.E. and Kosenko, A. (2024b). *The Economics of Information in a World of Disinformation: A Survey. Part 2: Direct Communication*. NBER Working Paper No. 32050. Cambridge, MA: National Bureau of Economic Research.

Vosoughi, S., Roy, D. and Aral, S. (2018). The Spread of True and False News Online. *Science*, 359(6380), pp. 1146–1151.

Weber, M. (1978). *Economy and Society*. Edited by G. Roth and C. Wittich. Berkeley: University of California Press.

Williamson, O.E. (1985). *The Economic Institutions of Capitalism: Firms, Markets, Relational Contracting*. New York: Free Press.

Zuboff, S. (2019). *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power*. New York: PublicAffairs.