



Munich Personal RePEc Archive

Does the Rise of Anthropogenic Signal a Sustainable AI Economy or Another Technological Bubble?

Ngunza Maniata, Kevin

University of Kinshasa

18 October 2025

Online at <https://mpra.ub.uni-muenchen.de/126512/>
MPRA Paper No. 126512, posted 27 Oct 2025 08:48 UTC

Does the Rise of Anthropic Signal a Sustainable AI Economy or Another Technological Bubble?

Kevin Ngunza Maniata*

Version 1.0 – October 2025

Abstract

The recent ascent of Anthropic, a United States-based artificial-intelligence company founded in 2021 by former OpenAI executives, has reignited the debate over whether the global AI boom represents sustainable technological transformation or a new financial bubble. With a private valuation surpassing 180 billion USD and projected annualized revenue exceeding 20 billion USD by 2026, Anthropic embodies both the promise of rapid innovation and the risks of speculative exuberance. This paper examines the firm's growth within the theoretical frameworks of Schumpeterian innovation, Minskyan financial cycles, and contemporary analyses of digital-economy concentration. Drawing on publicly available financial data, corporate disclosures, and secondary literature, it interprets Anthropic's trajectory as a case study in the financialization of cognition. The discussion highlights how alliances with Amazon and Google have turned frontier AI into an infrastructure-dependent oligopoly, while unresolved issues of data ownership and legal accountability question the durability of such valuations. The study concludes that Anthropic's rise illustrates the dual nature of modern technological capitalism: the capacity for exponential value creation tempered by systemic fragility and institutional lag.

Keywords: Artificial Intelligence; Anthropic; Financialization; Valuation; Technological Innovation; Industrial Organization; Intellectual Property; Speculative Cycles.

JEL Codes: O33, L86, G32, K11, M21.

Disclaimer: The views expressed in this paper are those of the author and do not necessarily reflect the positions of any institution. All data are drawn from public sources and academic literature. The paper is intended for scholarly discussion and does not constitute financial or policy advice.

* Department of Economics, University of Kinshasa, Democratic Republic of Congo. ORCID: <https://orcid.org/0000-0002-0758-6839>. Email: ngunza_kevin@yahoo.fr

1. Introduction

The dramatic rise of Anthropic has become one of the most striking financial and technological events of the mid-2020s. Founded in 2021 by former OpenAI researchers, including Dario and Daniela Amodei, the company has rapidly evolved from a research start-up to a firm valued at more than 180 billion USD following a 13-billion-dollar funding round in 2025. Its partnership with Amazon and Google, its family of language models known as Claude, and its projected annualized revenue above 20 billion USD by 2026 have led observers to ask whether this expansion represents a sustainable technological transformation or a new speculative episode similar to the Internet bubble at the turn of the century. As Robert Shiller (2000) reminded economists, “irrational exuberance” often accompanies technological change when expectations about future productivity outpace the evidence of current performance. The question is therefore not only empirical but conceptual: does the valuation of Anthropic reflect genuine innovation or the financialization of imagination?

Technological optimism has long been embedded in capitalist development. Schumpeter (1934) argued that innovation drives economic growth through “creative destruction,” a process in which new technologies disrupt established industries and generate temporary monopolies. Yet, as Minsky (1986) warned, financial systems tend to amplify such waves of innovation into speculative cycles when credit and capital concentrate around expected rather than realized returns. The current AI boom sits at this intersection of creativity and speculation. Firms such as Anthropic and OpenAI have transformed machine learning into a general-purpose technology that promises to reshape labor, production, and governance. At the same time, their valuations have reached levels difficult to justify through conventional profitability metrics, inviting comparison with historical episodes in which financial enthusiasm preceded market correction.

The purpose of this paper is to analyze Anthropic’s trajectory through the combined lenses of innovation economics, industrial organization, and financial theory. It seeks to determine whether the company’s rapid expansion indicates a self-sustaining technological ecosystem or the formation of a fragile bubble driven by narrative and capital flows. This inquiry contributes to the growing literature on the political economy of artificial intelligence by linking firm-level valuation to systemic structures of investment, regulation, and public accountability. While previous studies have focused on the technical performance of AI models or the social consequences of automation, fewer have examined the financial logic

that underpins the industry's consolidation. By integrating theoretical insights from Mazzucato (2013), Stiglitz (2021), and Perez (2002), the analysis situates Anthropic within a broader debate on how innovation, finance, and governance interact in late capitalism.

The significance of the question extends beyond one company. Artificial intelligence has become a strategic resource comparable to energy in the industrial era. Its concentration within a handful of private actors raises profound issues of economic sovereignty, regulatory capacity, and distributive justice. Whether this concentration represents an efficient allocation of resources or a symptom of market failure will shape not only the future of technology but also the evolution of global economic order. As Carlota Perez (2002) observed, technological revolutions often begin with speculative excess that finances infrastructure before policy and institutions catch up. Determining whether the AI surge has entered such a phase requires examining both financial indicators and the institutional context that sustains them.

This introduction frames the central research question: **Does Anthropic's rise represent sustainable value creation or a speculative distortion of market expectations?** The following analysis proceeds in three stages. First, it reviews the theoretical literature on innovation cycles, market concentration, and financial reflexivity. Second, it examines empirical evidence on Anthropic's valuation, revenue structure, and legal controversies. Finally, it discusses the policy and theoretical implications of these patterns for competition, intellectual-property governance, and macro-financial stability. The objective is not to predict short-term market movements but to clarify the economic logic that defines artificial intelligence as both an engine of growth and a potential source of systemic risk.

2. Methodological and Analytical Approach

This study adopts an interpretive and interdisciplinary analytical framework grounded in the traditions of political economy, innovation economics, and financial theory. The methodology is qualitative but theory-driven, combining descriptive data analysis with conceptual synthesis. The empirical material consists of publicly available financial information from press releases, corporate filings, and specialized media reports concerning Anthropic's funding, valuation, and product development between 2021 and 2025. These data are complemented by secondary academic literature on technology cycles, firm valuation, and market concentration. The purpose is not to model financial returns mathematically but to situate observed corporate behavior within established theoretical paradigms.

The analytical approach is built on three complementary pillars. The first is the Schumpeterian model of innovation that interprets technological change as a dynamic process of creative destruction (Schumpeter 1934). In this view, temporary monopolies and rising valuations are not anomalies but integral features of capitalist transformation. The second pillar is Minsky's theory of financial instability, which posits that periods of optimism and credit expansion can turn innovation into speculative excess when expectations of future profitability detach from underlying cash flows (Minsky 1986). The third pillar is the institutional perspective on financialization articulated by Mazzucato (2013) and Perez (2002), which emphasizes that innovation is socially and politically embedded. These frameworks allow the paper to interpret Anthropic's valuation both as an indicator of technological promise and as a reflection of systemic market behavior.

Data interpretation follows an inductive logic. Observed trends in Anthropic's capital inflows, revenue projections, and legal settlements are treated as empirical manifestations of deeper structural forces. The analysis draws on qualitative content from official announcements by Anthropic, Amazon, and Google, as well as coverage by credible technology and economic media outlets. Because precise revenue and cost data for private firms remain undisclosed, this study triangulates information from multiple independent sources to estimate approximate financial magnitudes. This methodological choice acknowledges the opacity of the AI industry while maintaining analytical rigor through cross-verification. The approach aligns with the interpretive-institutionalist tradition that regards economic phenomena as narratives mediated by power, expectation, and institutional context (MacKenzie 2006; Stiglitz 2021).

Quotations and citations are employed to anchor the discussion in the scholarly literature and to ensure transparency of argument. For instance, Shiller's (2000) observation that "speculative bubbles are natural expressions of human enthusiasm for progress" provides a conceptual bridge between financial exuberance and technological optimism. Likewise, Mazzucato's (2013) claim that the state is often "the first risk-taker and the last investor" informs the discussion of public accountability in AI funding. These references are not decorative but serve to position the present analysis within ongoing debates on how capital, knowledge, and regulation coevolve.

The analytical method also incorporates comparative reasoning. Anthropic is examined relative to other AI leaders such as OpenAI and DeepMind, allowing inferences about

structural patterns in the broader market. Comparative analysis strengthens external validity by distinguishing company-specific features from sector-wide dynamics. The evaluation of whether Anthropic's rise signals sustainable innovation or a speculative bubble therefore rests on three interrelated criteria: the proportionality of valuation to revenue potential, the degree of dependency on incumbent infrastructures, and the alignment between private incentives and social outcomes.

Finally, the methodology acknowledges its limitations. The rapid evolution of artificial intelligence makes any assessment provisional, and private-market data are inherently incomplete. The objective is interpretive clarity rather than predictive precision. By combining theoretical depth with empirical observation, the study seeks to clarify how financial narratives, institutional structures, and technological trajectories interact to produce what might be called the political economy of artificial cognition.

3. Literature Review and Analytical Framework

Economic scholarship on technological innovation and financial cycles has long recognized that major technological transformations tend to unfold through alternating phases of discovery, enthusiasm, and institutional stabilization. Joseph Schumpeter (1934) described this process as creative destruction, the dynamic by which new combinations of knowledge and capital disrupt existing structures and create temporary monopolies. In his view, capitalist progress depends on the continuous replacement of old production systems by new ones, yet each wave of innovation produces not only growth but also instability. Schumpeter's followers later observed that the profits of pioneering firms tend to be short lived unless they are reinforced by institutional or infrastructural advantages. In the contemporary artificial intelligence industry, these advantages take the form of access to data, computational resources, and financial capital.

Theories of financial instability add another dimension to this dynamic. Hyman Minsky (1986) argued that capitalist economies generate endogenous cycles of optimism and fragility. During periods of rapid innovation, credit expands as investors project high future earnings, leading to speculative valuations that may become disconnected from actual productivity. When expectations correct, asset prices fall and the cycle resets. Minsky's analytical framework explains why every major technological revolution, from the railway boom of the nineteenth century to the dot com bubble, has exhibited a pattern of exuberant investment followed by adjustment. Applied to artificial intelligence, his theory suggests that

the current wave of funding and valuation may represent the financial stage of a longer innovation cycle rather than purely rational market assessment.

Building on these classical perspectives, contemporary scholars such as Mariana Mazzucato (2013) and Carlota Perez (2002) have reframed innovation as a coevolutionary process between technology, finance, and institutions. Mazzucato emphasized that innovation is “not a spontaneous market outcome but the product of mission oriented public investment.” She demonstrated that many foundational technologies of the digital era emerged from state funded research before being commercialized by private firms. Perez extended this argument historically, showing that each technological revolution passes through two broad phases: an installation period driven by speculative finance and a deployment period in which institutions adapt and stabilize returns. According to her model, speculative bubbles can play a productive role by channeling capital into new infrastructures, but only if regulation and governance eventually align incentives with long term societal value. These frameworks provide a lens for analyzing Anthropic’s trajectory as part of an unfolding installation period of the artificial intelligence revolution.

Research in industrial organization and digital economics further clarifies how innovation in networked markets tends to produce concentration rather than dispersion. Shapiro and Varian (1999) demonstrated that information goods display strong network effects and high switching costs, which naturally lead to winner take all outcomes. Subsequent studies by Brynjolfsson and McAfee (2014) showed that digital platforms accumulate learning advantages through data feedback loops: each additional user or transaction improves the algorithmic product, reinforcing the incumbent’s dominance. The economics of scaling laws in machine learning described by Cowen and Tabarrok (2022) confirms that the performance of large models increases superlinearly with computing resources, further privileging firms that control massive infrastructure. This literature helps explain why the artificial intelligence sector is consolidating around a few entities such as OpenAI, Anthropic, DeepMind, and xAI, whose ability to finance and operate large models creates a structural barrier to entry.

At the intersection of finance and technology, Robert Shiller (2000) analyzed how narratives of progress sustain speculative booms. He observed that “economic bubbles are, in essence, social epidemics of excitement.” Market participants construct collective stories that justify extreme valuations, and these stories acquire credibility through repetition and media amplification. Donald MacKenzie (2006) later described such dynamics as performative,

meaning that financial models and expectations can shape the very markets they are meant to describe. In the context of artificial intelligence, discourses of frontier safety, alignment, and general intelligence function not only as scientific goals but also as narrative assets that attract investment. The financialization of artificial cognition thus operates through belief systems as much as through technical achievement.

The legal and ethical literature introduces a further layer of complexity. Intellectual property scholars such as Suzanne Scotchmer (2004) and William Nordhaus (1969) analyzed the tradeoff between protecting innovators and maintaining social access to knowledge. Excessive protection can hinder cumulative innovation, while insufficient protection undermines incentives to invest. The emergence of generative AI models that train on copyrighted data has disrupted this equilibrium. Recent cases of litigation and settlement, including Anthropic's 1.5-billion-dollar agreement to resolve copyright claims in 2025, illustrate how legal uncertainty can impose retrospective costs on innovation. Cohen and Reidenberg (2024) argue that data governance has become "the new terrain of property," transforming collective cultural resources into proprietary inputs. The theoretical challenge is to design legal frameworks that recognize the collective origins of data without stifling progress.

Political economy analyses of digital capitalism, such as those of Zuboff (2019) and Stiglitz (2021), highlight that data concentration and algorithmic control reproduce forms of market power analogous to industrial monopolies. Zuboff's concept of surveillance capitalism describes how behavioral data become raw material for profit extraction. In the case of Anthropic and its peers, the commodified asset is no longer behavior but cognition itself, the predictive and generative capacity of models that approximate reasoning. Stiglitz emphasized that when markets fail to internalize externalities, private incentives diverge from social welfare, producing outcomes that are both inefficient and inequitable. Applying his argument to artificial intelligence suggests that the race for scale may generate systemic risk, as the benefits of technological capability accrue to a narrow set of firms while costs such as energy consumption, legal disputes, and labor displacement are socialized.

Taken together, these bodies of literature establish the conceptual foundation for the present analysis. The Schumpeterian emphasis on innovation, the Minskyan account of financial instability, and the institutional theories of Mazzucato and Perez provide complementary lenses through which Anthropic's rise can be understood. The economics of digital platforms

and intellectual property governance further reveal how structural forces of concentration shape the distribution of technological gains. Finally, sociological and ethical perspectives underscore the role of narratives, law, and power in determining whether technological revolutions culminate in sustainable growth or in speculative correction. Within this framework, Anthropic's valuation surge becomes not merely a financial event but a revealing case study in the evolving political economy of artificial intelligence.

4. Empirical Context and Analytical Discussion

Anthropic's trajectory from its founding in 2021 to its 2025 valuation epitomizes the financial acceleration of artificial intelligence. Established by Dario and Daniela Amodei after their departure from OpenAI, Anthropic began with a mission to develop safe and interpretable large language models. Within four years, it attracted cumulative investments exceeding 13 billion USD and achieved a private valuation of approximately 183 billion USD. Amazon committed up to 4 billion USD in late 2023, and Google contributed an additional 2 billion USD in 2024, making these two corporations both investors and infrastructure providers. According to company reports and industry analysts, Anthropic's annualized revenue reached nearly 7 billion USD in the third quarter of 2025, with internal projections targeting 20 to 26 billion USD in 2026. The implied revenue to valuation multiple exceeds 1 to 25, a ratio reminiscent of high growth technology valuations during earlier speculative periods such as the late 1990s Internet bubble (Shiller 2000).

The company's strategic positioning differentiates it from OpenAI, its principal competitor. While OpenAI has cultivated a large consumer base through ChatGPT subscriptions, Anthropic has concentrated on enterprise clients. Its Claude family of models, including Claude 3, Claude Opus, and the compact Claude Haiku 4.5, has been integrated into business applications for document processing, compliance analysis, and customer service. This enterprise orientation produces relatively stable recurring revenue, as corporate clients sign multi year contracts and incur high switching costs. The economics of this model mirror those of software as a service firms but with higher infrastructure dependency. Each training cycle requires extensive computational capacity, primarily supplied by Amazon Web Services and Google Cloud. Consequently, Anthropic's growth strategy reinforces the oligopolistic structure of the cloud industry, transforming technological innovation into a function of existing corporate ecosystems rather than independent competition (Brynjolfsson and McAfee 2014).

Anthropic’s financial and strategic configuration can be summarized by the following indicators:

Year	Estimated Revenue (USD Billion)	Cumulative Investment (USD Billion)	Approximate Valuation (USD Billion)	Principal Investors
2021	<0.1	0.7	4	Early-stage venture funds
2022	0.5	2	15	Google (initial partnership)
2023	3	6	80	Amazon (up to 4.0 commitment)
2024	5	9	150	Google (additional 2.0)
2025	7	13	183	Amazon, Google, major venture capital funds

Sources: company disclosures, media reports, and market analyses compiled from public sources, 2021–2025.

This empirical profile illustrates the interplay of innovation and speculation described in the theoretical literature. The rapid escalation of valuation relative to revenue suggests that investors are capitalizing expectations of future dominance rather than present profitability. As Minsky (1986) explained, “stability itself can be destabilizing” when optimism leads to leveraged expansion. In Anthropic’s case, optimism rests on assumptions that scaling laws will continue to yield exponential performance gains and that enterprise demand for AI services will remain price inelastic. Such assumptions may be fragile if technological progress slows or if clients face integration limits. The historical analogy to the semiconductor and Internet sectors indicates that productivity gains often materialize more slowly than financial markets anticipate (Perez 2002).

Legal and ethical challenges complicate the financial narrative. In September 2025, Anthropic agreed to pay 1.5 billion USD to settle a class-action lawsuit alleging copyright infringement in its training data. The plaintiffs claimed that the company’s models had used protected works without authorization. The settlement, described by legal analysts as the largest publicized intellectual-property recovery in the AI industry, highlighted the unresolved question of data ownership. Economically, the payment can be interpreted as a retroactive licensing cost, effectively converting legal risk into an ex-post expenditure. For firms with

large capital reserves, such costs are absorbable; for smaller competitors, they are existential. This asymmetry reinforces market concentration by allowing well-capitalized firms to internalize legal uncertainty while deterring new entrants (Scotchmer 2004). As Cohen and Reidenberg (2024) observed, “the privatization of training data transforms shared cultural resources into corporate assets,” an observation that encapsulates the core tension between innovation and equity in generative AI.

The relationship between Anthropic and its corporate investors exemplifies a form of structural dependency that blurs the boundary between competition and vertical integration. Amazon and Google are simultaneously investors, suppliers, and customers. Their cloud platforms provide the computing infrastructure for model training, while their enterprise networks channel clients toward Anthropic’s products. This triangular configuration recalls the industrial arrangements described by Chandler (1990) in which capital-intensive technologies produce economies of scale that deter independent competition. The result is a hybrid structure that combines the agility of a start-up with the systemic importance of a conglomerate. Such interdependence complicates regulatory oversight because no single entity controls the entire value chain, yet the combination of capital and infrastructure effectively centralizes power.

Macroeconomic indicators provide additional context. Global venture funding in AI reached more than 70 billion USD in 2025, accounting for nearly a quarter of all technology investment worldwide. At the same time, productivity statistics in advanced economies show only modest improvements. This divergence between financial expansion and real productivity growth mirrors the phenomenon of financialization described by Epstein (2005), in which capital accumulation becomes increasingly detached from tangible output. Anthropic’s valuation therefore participates in a broader pattern of capital seeking returns in expectation rather than realization. As Stiglitz (2021) argued, such dynamics can lead to misallocation of resources when private valuation does not correspond to social value.

Supporters of the AI boom counter that these valuations reflect the unprecedented potential of general-purpose technologies. They argue that just as electricity or computing required decades to diffuse before transforming productivity, artificial intelligence will yield long-term gains once infrastructure, skills, and regulation adapt. Proponents view high valuations as rational bets on the future of intelligence automation rather than speculative distortions. They note that firms such as Anthropic invest heavily in safety research and model interpretability,

which may accelerate institutional adoption. From this perspective, market optimism functions as a financing mechanism for long-term innovation, not as a symptom of mania. Perez (2002) provided a historical precedent for this argument when she wrote that “every bubble leaves behind the fixed capital that becomes the foundation for the next golden age.” If AI valuations finance data centers, research talent, and algorithmic breakthroughs, the eventual social returns could justify present exuberance.

Yet, even within this optimistic interpretation, sustainability depends on institutional adaptation. The history of technological revolutions shows that the transition from speculative installation to stable deployment requires policy innovation as well as technical progress. In the absence of clear regulation, excessive concentration or data appropriation could provoke political backlash and erode public trust. The challenge for regulators and central banks is to determine whether the AI sector constitutes a source of systemic risk comparable to finance or energy. The Financial Stability Board (2011) has already suggested that large technology firms may exhibit “systemic importance” due to their control of digital infrastructure. If cognitive infrastructure becomes as indispensable as payment systems, similar oversight mechanisms may eventually be required.

Anthropic’s projected tripling of annual revenue thus represents both an economic milestone and a conceptual stress test for existing theories of innovation and valuation. The firm’s success demonstrates the scalability of machine intelligence, yet its structure and valuation raise the question of whether value creation is proportional to productive contribution. The empirical evidence points to a dual reality: genuine technological progress coexists with financial speculation. Determining which element dominates will depend on whether the institutional environment can translate private incentives into sustainable public value.

5. Policy and Theoretical Implications

The evidence surrounding Anthropic’s rise calls for a reexamination of how economic institutions conceptualize innovation, competition, and financial stability in the age of artificial intelligence. Traditional industrial organization models assume that market concentration is a temporary consequence of technological leadership. However, the artificial intelligence sector exhibits structural features that may entrench dominance rather than erode it. The combination of data network effects, scale economies in computing power, and high capital requirements suggests that concentration is an inherent characteristic of the market. Shapiro and Varian (1999) observed that “the economics of information goods naturally favor

winner who can capture and hold markets.” Anthropic’s alliance with Amazon and Google exemplifies this tendency, as the firm’s dependence on their cloud infrastructure and capital simultaneously reinforces and legitimizes their market power. In this environment, competition policy must evolve beyond static antitrust principles toward dynamic regulation that considers access to computational and data infrastructures as a determinant of market power.

Current competition frameworks focus on preventing collusion or predatory pricing, yet these instruments are poorly suited to address networked interdependencies. Regulators may need to treat computational resources and model architectures as essential facilities, requiring transparency and fair access to ensure pluralism. The European Union’s Digital Markets Act represents an early attempt to redefine gatekeeping in the digital economy, but its focus remains on consumer platforms rather than industrial artificial intelligence ecosystems. As Cowen and Tabarrok (2022) suggested, the economics of scaling laws in artificial intelligence may create a new class of natural oligopolies where efficiency and concentration coincide. The challenge is to distinguish acceptable efficiency from strategic foreclosure. Without such differentiation, market dominance could solidify into a form of cognitive monopoly that limits innovation rather than enabling it.

Intellectual property law represents another critical domain of policy adjustment. The legal disputes surrounding training data underscore the inadequacy of existing frameworks to manage machine learning’s dependence on cultural and informational commons. Nordhaus (1969) and Scotchmer (2004) demonstrated that the optimal design of intellectual property rights balances incentives for creators with the diffusion of knowledge. Generative artificial intelligence models blur this balance because their learning process transforms public information into private capability. Cohen and Reidenberg (2024) described this as “the enclosure of data,” a modern parallel to the privatization of shared resources during the industrial revolution. Policymakers could mitigate this by establishing collective licensing mechanisms or data trusts that compensate content creators while maintaining access for innovation. Such structures would align with Mazzucato’s (2013) argument that public institutions must reclaim an entrepreneurial role in directing technological development toward inclusive outcomes.

Macroeconomic and financial stability considerations are equally important. The scale of capital inflows into artificial intelligence has introduced a new asset class whose valuations

are largely based on expectations rather than tangible output. Central banks and regulators traditionally monitor leverage and liquidity in housing or credit markets, yet they seldom consider the systemic effects of intangible asset bubbles. Perez (2002) and Minsky (1986) showed that financial cycles linked to technological revolutions often trigger resource misallocation when enthusiasm outruns absorptive capacity. A sudden correction in artificial intelligence valuations could reverberate through venture capital portfolios, pension funds, and technology linked exchange traded funds. Incorporating technological sector indicators into macroprudential stress tests would therefore be prudent. As Stiglitz (2021) warned, when markets fail to internalize systemic risk, public authorities bear the burden of stabilization.

At the same time, the artificial intelligence boom presents an opportunity to rethink fiscal and industrial policy. Because large scale artificial intelligence models rely on infrastructures built through public investment such as electricity grids, data networks, and education systems, governments have a legitimate claim to shape their development. Mazzucato (2013) emphasized that the state often acts as “the first risk taker,” yet rarely captures proportional returns. One corrective measure would be to implement public co investment models or taxation schemes that recycle part of the profits from algorithmic capital into public goods such as education, research, and environmental sustainability. As Atkinson and Stiglitz (1980) argued, progressive taxation under monopolistic competition can enhance both equity and efficiency by redistributing rents derived from market power. Applying this logic to artificial intelligence could ensure that the productivity gains of machine cognition do not concentrate excessively among a few global firms.

Ethical and governance frameworks also require strengthening. Anthropic presents itself as a public benefit corporation committed to safety and alignment. While this status allows directors to balance social and shareholder objectives, research on corporate governance suggests that profit motives remain dominant in high growth sectors (Hart and Zingales 2017). The tension between ethical commitments and financial imperatives may become acute as competition intensifies. Russell (2021) observed that ensuring “human compatible” artificial intelligence depends on external accountability rather than self regulation. Public agencies should therefore establish independent mechanisms for auditing model behavior, verifying training data, and assessing environmental and social impacts. These institutions could mirror the role of financial auditors, providing assurance that claims of safety and fairness correspond to verifiable standards.

The global dimension of artificial intelligence development further complicates the policy landscape. Advanced economies dominate both capital investment and computational infrastructure, while emerging economies risk becoming passive consumers of proprietary cognitive services. This asymmetry mirrors what Prebisch (1950) described as the structural dependence of peripheral economies on industrial centers. In the digital era, data replace raw materials as the resource extracted from the periphery, and trained models replace manufactured goods as the high value imports. To counter this pattern, international organizations and regional development banks could support open source artificial intelligence projects, multilingual data initiatives, and regional compute facilities. Such strategies would democratize access to machine learning and reduce cognitive dependency. Empirical evidence from digital infrastructure programs shows that localized innovation ecosystems enhance economic resilience and foster inclusive growth.

Beyond specific regulatory measures, Anthropic's rise invites theoretical reflection on the nature of value creation in a knowledge economy. Classical economic models conceptualize value as the product of labor and capital, while modern growth theory adds knowledge as a factor of production. The proliferation of artificial intelligence introduces a new form of capital: cognitive capital, defined as the capacity of algorithms to perform reasoning and generative tasks once limited to human intellect. This development challenges the boundaries of existing economic categories. When cognition itself becomes a commodity, the distinction between tool and producer collapses. As Cao & Meng (2025) observed in their survey of Chinese public attitudes, "we sorted AI ethics concerns into eight dimensions: safety, transparency, fairness, personal data protection, liability, truthfulness, human autonomy, and human dignity." Economists must therefore reconsider how to measure productivity, labor, and welfare in an environment where machines contribute to thought as well as production.

From a broader theoretical standpoint, Anthropic's ascent exemplifies the convergence of innovation and finance into a unified system of speculative knowledge production. Financial markets no longer merely fund innovation; they co produce it by shaping the expectations that guide research and deployment. MacKenzie (2006) described this performative relationship as a feedback loop between valuation and invention. The political economy of artificial intelligence thus transcends the boundaries of both economics and technology. It reveals how belief, narrative, and power operate as productive forces. Whether the artificial intelligence revolution becomes a foundation for long term prosperity or a transient bubble will depend on the capacity of institutions to align these forces with collective goals.

Conclusion

Anthropic's extraordinary ascent captures a historical turning point in the evolution of technological capitalism. The company's rise from a research initiative in 2021 to a firm valued above one-hundred-eighty billion USD by 2025 reflects both genuine innovation and speculative intensity. Its partnership structure with Amazon and Google, its rapid revenue expansion, and its costly legal settlements all reveal the dual nature of frontier artificial intelligence: a field that generates enormous productive potential while simultaneously magnifying systemic fragilities. The analysis developed in this paper demonstrates that Anthropic's valuation cannot be interpreted solely as a signal of technological progress nor dismissed simply as financial excess. It represents a complex equilibrium between capital, expectation, and institutional adaptation.

From the standpoint of economic theory, the Anthropic case validates Schumpeter's vision of creative destruction but also confirms Minsky's warning that prosperity often breeds instability. The firm's growth illustrates how speculative finance can accelerate innovation by funding large-scale experimentation, yet it also shows that unchecked concentration can transform technological revolutions into mechanisms of rent extraction. Whether the present wave of investment produces sustainable growth or a corrective downturn will depend on how effectively regulatory institutions, fiscal policies, and international cooperation channel private incentives toward public welfare.

Policy makers face the task of ensuring that artificial intelligence evolves as a productive infrastructure rather than a speculative arena. This involves redefining competition policy to include computational access, reforming intellectual-property law to balance innovation and fairness, and developing macro-prudential tools to monitor intangible-asset cycles.

Governments should also consider mechanisms for reinvesting algorithmic rents into education and digital inclusion so that cognitive automation enhances rather than replaces human capability. In emerging economies, open-source collaboration and regional compute initiatives could mitigate dependency on global oligopolies.

Ultimately, the question of whether Anthropic's rise signals a sustainable AI economy or another technological bubble cannot yet be resolved empirically. What can be asserted is that the sustainability of this new economy will depend less on the ingenuity of algorithms than on the wisdom of governance. The future of artificial intelligence will be determined by whether financial capital continues to dominate cognitive capital or whether societies learn to

align both within a framework of accountability and shared prosperity. The outcome of that alignment will define not only the trajectory of Anthropoc but the architecture of twenty-first-century capitalism itself.

References

- Arthur, W. B. (1989). *Competing technologies, increasing returns, and lock-in by historical events*. *Economic Journal*, 99 (394), 116-131. <https://doi.org/10.2307/2234208>
- Atkinson, A. B., & Stiglitz, J. E. (1980). *Lectures on Public Economics*. New York, NY: McGraw-Hill.
- Boltanski, L., & Chiapello, È. (2007). *The New Spirit of Capitalism*. London, UK: Verso.
- Brynjolfsson, E., & McAfee, A. (2014). *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*. New York, NY: W. W. Norton.
- Chandler, A. D. (1990). *Scale and Scope: The Dynamics of Industrial Capitalism*. Cambridge, MA: Harvard University Press.
- Christensen, C. M. (1997). *The Innovator's Dilemma: When New Technologies Cause Great Firms to Fail*. Boston, MA: Harvard Business School Press.
- Cowen, T., & Tabarrok, A. (2023). *How to Learn and Teach Economics with Large Language Models, Including GPT*. GMU Working Paper in Economics No. 23-18. Available at SSRN: <https://doi.org/10.2139/ssrn.4391863> (SSRN)
- Damodaran, A. (2023). *AI's Winners, Losers and Wannabes*. New York University, Stern School of Business. Available at: <https://pages.stern.nyu.edu/~adamodar/pdffiles/country/AI.pdf> (Stern School of Business)
- Epstein, G. A. (2005). *Financialization and the World Economy*. Cheltenham, UK: Edward Elgar.
- Hart, O., & Zingales, L. (2017). Companies should maximize shareholder welfare — not market value. *Journal of Law, Finance & Accounting*, 2(2), 247-274. <https://doi.org/10.1561/108.00000022>
- MacKenzie, D. (2006). *An Engine, Not a Camera: How Financial Models Shape Markets*. Cambridge, MA: MIT Press.

- Mazzucato, M. (2013). *The Entrepreneurial State: Debunking Public vs. Private Sector Myths*. London, UK: Anthem Press.
- Minsky, H. P. (1986). *Stabilizing an Unstable Economy*. New Haven, CT: Yale University Press.
- Nordhaus, W. D. (1969). *Invention, Growth, and Welfare: A Theoretical Treatment of Technological Change*. Cambridge, MA: MIT Press.
- Perez, C. (2002). *Technological Revolutions and Financial Capital: The Dynamics of Bubbles and Golden Ages*. Cheltenham, UK: Edward Elgar.
- Prebisch, R. (1950). *The Economic Development of Latin America and Its Principal Problems*. New York, NY: United Nations.
- Russell, S. J. (2021). *Human Compatible: Artificial Intelligence and the Problem of Control*. New York, NY: Viking.
- Schiller, R. J. (2000). *Irrational Exuberance*. Princeton, NJ: Princeton University Press.
- Schumpeter, J. A. (1934). *The Theory of Economic Development: An Inquiry into Profits, Capital, Credit, Interest, and the Business Cycle*. Cambridge, MA: Harvard University Press.
- Scotchmer, S. (2004). *Innovation and Incentives*. Cambridge, MA: MIT Press.
- Shapiro, C., & Varian, H. R. (1999). *Information Rules: A Strategic Guide to the Network Economy*. Boston, MA: Harvard Business School Press.
- Stiglitz, J. E. (2021). Markets, power, and inequality: The political economy of our time. *American Economic Review*, 111(5), 1588-1627. <https://doi.org/10.1257/aer.111.5.1588>
- Zuboff, S. (2019). *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power*. New York, NY: PublicAffairs.