


Agentic Inequality



Authors: Matthew Sharp, Omer Bilgin, Iason Gabriel, Lewis Hammond

Agentic Inequality

Matthew Sharp^{1,2} Omer Bilgin³ Iason Gabriel⁴ Lewis Hammond^{3,5}

24 April 2026

¹Oxford Martin AI Governance Initiative ²Access Partnership ³University of Oxford

⁴Google DeepMind ⁵Cooperative AI Foundation

Abstract

Autonomous AI agents capable of complex planning and action mark a shift beyond today’s generative tools. As these systems enter political and economic life, who can access them, how capable they are, and how many can be deployed will shape distributions of power and opportunity. We define this emerging challenge as “agentic inequality”: disparities in power, opportunity, and outcomes arising from unequal access to, and capabilities of, AI agents. We show that agents could either deepen existing divides or, under the right conditions, mitigate them. The paper makes three contributions. First, it develops a framework for analysing agentic inequality across three dimensions: *availability*, *quality*, and *quantity*. Second, it argues that agentic inequality differs from earlier technological divides because agents function as autonomous delegates rather than tools, generating new asymmetries through scalable goal delegation and direct agent-to-agent competition. Third, it analyses the technical and socioeconomic drivers likely to shape the distribution of agentic power, from model release strategies to market incentives, and concludes with a research agenda for governance.

1. Introduction

Generative AI models – typified by systems that convert prompts into prose, code or images – now underpin AI agents: systems that pursue complex goals with significant autonomy by perceiving their environment, planning, and executing multi-step tasks across digital environments (Wang et al., 2024; Kasirzadeh and Gabriel, 2025). These systems vary in autonomy, efficacy, goal complexity, and generality, ranging from narrow task executors operating within constrained workflows to general-purpose delegates that manage cross-domain goals. Although still at an early stage, AI agents are beginning to perform economically valuable workflows, and ongoing research aims to extend their capabilities. This shift from AI as a command-following tool to AI as an autonomous goal-pursuing actor matters because it affects human agency: the capacity of individuals to pursue goals that they value (Sen, 1999; Prunkl, 2024).

This paper introduces and explores “agentic inequality”: disparities in power and outcomes that may arise from unequal access to AI agents. We argue that agentic inequality differs from earlier technological divides for two reasons. First, agents create new power asymmetries by allowing users to delegate complex tasks at scale. Second, they create new competitive dynamics through direct agent-to-agent competition. Our argument sits within broader debates about the political

economy of AI, but focuses specifically on disparities arising from the distinctive features of AI agents. As these systems are not yet widely deployed, a window of opportunity exists to proactively shape their development and ensure their benefits are distributed equitably.

We build on scholarship on the co-evolution of technology and society (Mumford, 1967; Jasanoff, 2004), and extend research on LLM impacts (Acemoglu and Restrepo, 2020; Brynjolfsson et al., 2025b; Noy and Zhang, 2023) and the “digital divide” (Selwyn, 2004; van Dijk, 2005; Warschauer, 2003). On that basis, we develop an analytical framework for this emerging challenge, exploring both the risks of new forms of inequality and the potential for agents to mitigate existing divides. Agents could amplify the advantages of well-resourced actors, but if capable systems are widely accessible, they could also reduce existing disadvantages by simplifying complex processes and performing tasks that previously required costly expertise. A critical step involves shifting our evaluative lens to view AI agents not just as individual products, but as a form of “infrastructure” that is likely to mediate access to essential goods, services, and opportunities (Plantin et al., 2018; Rothschild et al., 2025; Bergman et al., 2024; Chan et al., 2025; Lazar, 2025). Our focus is primarily on individuals and organisations, although the implications extend to all aspects of society.

The paper proceeds as follows. Section 2 identifies the key dimensions through which agentic inequality may manifest. Section 3 examines its implications across key societal domains. Section 4 analyses the technical and socio-political drivers that may widen or reduce it. Section 5 considers the governance challenge and sets out a research agenda.

2. Dimensions of Agentic Inequality

Agentic inequality has three dimensions: *availability*, *quality*, and *quantity*. Each captures a distinct and analytically separable source of differential advantage; together, they determine the extent to which an actor can derive value from agentic AI.

2.1. Availability

Availability is the most basic dimension. It has two layers. The first is *structural availability*: whether an agent is, in principle, available to a given actor, given market presence, regulatory permissibility, and the necessary infrastructure. The second is *accessibility*: whether that actor can make practical use of the agent, given constraints such as affordability, digital literacy, language support, and usability. Research on the digital divide has long emphasised this distinction, showing that formal access does not automatically translate into effective use (van Dijk, 2005; Selwyn, 2004; Warschauer, 2003). Together, these two conditions determine whether an individual or organisation can use even a single AI agent, which is the minimum requirement for participation in the agentic economy (Tomašev et al., 2025).

A further distinction is important here. Access to a foundation model is not the same as access to a functional agent. Foundation models are becoming more accessible, but, at the time of writing, deploying a capable agent remains difficult. The infrastructure needed to do so, including high-quality scaffolding, is still new, complex, and often concentrated among a small number of providers. Effective deployment also typically requires specialised expertise (Chan et al., 2025; Liu et al., 2025). Awareness and knowledge gaps therefore remain barriers to entry in their own right. Inequalities in availability may also reflect choice, where individuals or organisations decide not to use agents for ethical or practical reasons.

2.2. Quality

Availability alone does not settle what an agent enables its user to do. A second dimension is

quality. This refers to what an individual agent can do and the operational characteristics that determine its behaviour. Agent quality can vary along several axes:

- **Core intelligence:** the combination of an agent’s world knowledge, multimodal capabilities across language and other data types, and its ability to reason, plan, and solve problems. This depends largely on the sophistication of the underlying foundation model and its training data (Zhu et al., 2025).
- **Operational speed and throughput:** an agent’s efficiency depends on latency and data-processing capacity, which in turn reflect the computational resources supporting it, including hardware accelerators and cloud infrastructure (Sastry et al., 2024).
- **Reliability:** the consistency of an agent’s performance across tasks, including how robust it is and how often it fails.
- **Tool use:** an agent’s ability to access and effectively use external systems, such as APIs, proprietary databases, real-time data feeds, or physical actuators (Chan et al., 2025).
- **Disposition:** an agent’s behavioural tendencies, which may affect performance in context-specific ways. An aggressive disposition may help in negotiation, for example, but hinder performance in a customer service setting where politeness and helpfulness matter more.

These sub-components need not move together. Falling inference costs may commoditise core intelligence even as tool use remains gated by infrastructure controllers. As a result, *quality* gaps may narrow on one axis while remaining wide on another. From an inequality standpoint, *core intelligence* and *tool use* are likely to matter most. The former determines capacity for the complex, high-value tasks where agentic advantage is greatest; the latter determines whether an agent can act in the world rather than merely advise. *Disposition* may also become more important as multi-agent interaction expands, since behavioural tendencies in competitive or cooperative settings can directly shape outcomes.

2.3. Quantity

The final dimension is the *quantity* of agents an individual or organisation can deploy. This captures a source of advantage that is distinct from the capability of any single agent. Scale

can matter because running more agents allows a task to be pursued through multiple reasoning paths in parallel, rather than sequentially. Recent research finds that performance on complex tasks can improve as more agents are deployed, with larger gains on harder tasks (Li et al., 2024). Other work shows that running multi-agent teams in parallel can reduce latency while preserving accuracy, and in some cases improve task completion rates (Zhang et al., 2025). *Quantity* may therefore confer an advantage not just by adding effort, but by enabling broader search and faster convergence across complex problems.

2.4. Interaction Effects, User Dependence, and User-Dependent Value

These dimensions interact rather than operating independently. In some settings, *quantity* can partly offset *quality*: multiple competent but individually limited agents can outperform a single stronger agent operating alone, especially when they draw on diverse models, prompts, tools, or roles (Li et al., 2024; Żywot et al., 2026). In other settings, however, adding more agents yields little, as homogeneous systems produce increasingly correlated outputs and returns diminish sharply (Yang et al., 2026; Xu et al., 2026). Where quality, diversity, and coordination improve together, the gains can compound: heterogeneous multi-agent systems can outperform both single-agent baselines and larger but less differentiated agent pools (Yang et al., 2026; Żywot et al., 2026).

The value users derive from agents is also not fixed. Even where users have access to the same agent, the benefits they obtain may differ substantially. Outcomes depend not only on access but also on complementary assets such as workflow integration, managerial oversight, domain knowledge, and the capacity to deploy agents strategically (Weidinger et al., 2021; Bommasani et al., 2022). Some users may get much more out of the same agent than others. So even if access is equal in formal terms, outcomes may still diverge.

A related issue is *user dependence*: the extent to which people come to rely on agents for tasks they previously handled themselves. This dependence is unlikely to be evenly spread. People with fewer alternative resources may have stronger reasons to rely heavily on agents, which can leave them more exposed to quality problems or service interruptions (Prunkl, 2024). For that reason, dependence is better seen as a factor that intensifies the inequalities described above than as a separate fourth dimension. That said, dependence is

not inevitable: agents designed to explain their reasoning and build user competence over time could reduce reliance rather than deepen it.

Taken together, these dimensions make concrete two features that distinguish agentic inequality from earlier technological divides. The first is scalable delegation: deploying many agents in parallel may look less like using a better tool and more like directing a distributed workforce. The second is direct agent-to-agent competition: as agents negotiate, bid, and strategise on behalf of their users, the contest that previously took place between people may now turn on differences in algorithmic capability operating at machine speed (Zhu et al., 2025).

3. Implications of Agentic Inequality

Agentic inequality is not only a theoretical concern. It may also have substantial economic, social, and political effects. Whether agents mainly concentrate power among already well-resourced actors or broaden access to opportunity will depend on how differences in *availability*, *quality*, and *quantity* unfold across domains, and on the policy choices made in response. The discussion below draws on existing evidence where possible, projects from present trends where the technology is still emerging, and in some places considers longer-run possibilities. The text indicates which of these modes is being used.

3.1. Economic Impacts

3.1.1 Labour Market Effects

The labour market effects of AI remain uncertain. The available evidence is still limited, and most of it predates autonomous agents. “Agentic capital”, autonomous systems capable of managing entire business processes, could accelerate a shift in national income from labour to capital, extending a familiar dynamic from earlier capital-augmenting technologies (Brynjolfsson and McAfee, 2014; Brynjolfsson et al., 2025b). Yet the existing evidence on generative AI’s labour market impact is mixed. Economy-wide analyses suggest relative stability (Gimbel et al., 2025), while disaggregated studies of exposed industries find declining hiring for junior roles alongside stable or growing employment for senior workers (Brynjolfsson et al., 2025a; Hosseini Maasoum and Lichtinger, 2025). Some firm-level studies document a “levelling-up” effect, in which technology raises the productivity of less-experienced staff (Peng et al., 2023; Noy and Zhang, 2023; Cui et al., 2026; Brynjolfsson et al.,

2025b). That effect, however, is a feature of human-tool augmentation; it is far from clear that it would survive a transition to agentic capital built to execute and manage entire workflows with limited human input.

3.1.2 Industrial Organisation and Market Structures

In industry, the central question is whether agents intensify concentration or widen competition. Unequal deployment could accelerate the rise of “superstar firms” through two mechanisms. First, dominant firms may be able to use proprietary data to build more capable agents, improving services, attracting more users, and generating still more data in a feedback loop that smaller rivals may find difficult to match (Korinek and Vipra, 2025). Second, larger firms may be able to deploy greater numbers of agents to automate internal processes and accelerate innovation at a scale beyond the reach of smaller competitors. Yet the opposite dynamic is also possible. The declining cost and increasing availability of powerful agentic platforms could also act as an equalising force by lowering barriers to entry, enabling startups to run operations that previously required substantial human capital.

3.1.3 Consumer Welfare and Negotiations

Consumer markets may be among the first domains in which agentic inequality becomes visible. Corporate agents could amplify manipulative practices already documented in digital markets, including hidden information, scarcity cues, and deceptive choice architecture, at a scale and speed no human sales operation could match (Mathur et al., 2019; Kolt, 2025). This asymmetry may be especially pronounced in negotiations, where differences in algorithmic capability translate directly into worse outcomes for the weaker party (Zhu et al., 2025). At the same time, widely available consumer agents could pull in the opposite direction by automating comparison shopping, flagging deceptive pricing, and negotiating on users’ behalf at a scale no individual could manage alone (Van Loo, 2019; Zhu et al., 2025).

3.2. Social and Political Impacts

3.2.1 Access to Essential Services

Agents could reshape the citizen-state relationship, either producing a more stratified system or widening access to public services (Ilves et al., 2025). One risk is that affluent individuals may use premium agents to navigate complex bureaucratic workflows and optimise applications

in ways that secure better outcomes in health-care or legal systems (Bovens and Zouridis, 2002). That risk may grow if public services are designed around agentic interaction, leaving citizens without access to comparable tools at a disadvantage. But the reverse is also possible. “Public-good” agents could widen access by automating form-filling and proactively pursuing entitlements on behalf of citizens, reducing administrative burdens that fall disproportionately on lower-income groups. Taiwan’s experience with public digital infrastructure offers an instructive precedent (Weyl et al., 2024).

3.2.2 Political Discourse and Participation

In politics, the key tension is between amplifying elite influence and enabling broader participation (Summerfield et al., 2025). Differences in agent *quantity* and *quality* could allow well-resourced actors to deploy sophisticated agent swarms to run coordinated, multi-step influence campaigns, for example by overwhelming public consultations or generating personalised propaganda at scale (Gabriel et al., 2024; Persily and Tucker, 2020). At the same time, broad access to agents that can autonomously draft policy proposals, manage activist campaigns, and engage in political processes could strengthen the capacity of citizens and grassroots groups. Recent work has begun to explore how AI might support participatory deliberation in digital public spaces at scale (Goldberg et al., 2024). This shift from human persuasion to machine-mediated influence also raises a deeper question for democratic legitimacy, since the “public will” may come to be shaped more heavily by delegated political action.

3.2.3 Social Stratification

Agents may also introduce more explicit competition into social life, with effects that could run in either direction. They could deepen the “Matthew Effect”, whereby initial advantages compound over time, or they could widen social mobility. Unlike tools that merely augment human effort, agents could enable the scalable delegation of complex social strategies, from managing a professional’s networking to improving educational opportunities for one’s children. More broadly, they could reduce the transaction costs that inhibit beneficial coordination, enabling new forms of decentralised social bargaining (Krier, 2025). Scalable delegation of this kind could create a more persistent divide between those able to use such autonomous delegates effectively and

those who are not, as repeated small advantages widen gaps in social capital over time. But accessible agents could also improve mobility by offering coaching and by opening up opportunities that might otherwise remain out of reach.

4. Forces Shaping Agentic Inequality

The effects of AI agents will not be determined by the technology alone. They will emerge from a mix of technical constraints, market incentives, institutional arrangements, and political choices. Many of these forces apply to AI more broadly. They bear with particular force on agents, however, because systems that act autonomously on behalf of their users convert distributional disparities more directly into differences in outcomes. This preliminary mapping focuses on proximate technical and socioeconomic drivers; it does not attempt a full structural account. Factors such as affordability, access to capital, and digital literacy are themselves distributed unevenly along broader lines of race, class, and gender. Those wider patterns deserve further interdisciplinary attention.

4.1. Supply-Side and Ecosystem Drivers

4.1.1 Compute Costs and Capital Barriers

The political economy of AI is shaped in large part by the economics of computation. One result is a tension between concentrated capital and more distributed application. Two cost structures matter most. The first is the very large capital investment required to train a foundation model, which concentrates that capability within a small number of major technology firms (Sastry et al., 2024; Besiroglu et al., 2024). This affects both the availability of the most powerful agents and the level of core intelligence they can provide. The second is the ongoing cost of inference. Although much lower than training costs, inference can still become expensive as tasks grow more complex. More capable models also tend to command higher per-token prices because they require more compute, which creates a further pay-for-performance pattern in which premium tiers offer stronger reasoning or faster responses. At the same time, performance gains may taper off for many tasks. If so, falling inference costs could still make broadly accessible, “good enough” agents viable for baseline autonomous uses even if frontier systems retain a quality advantage.

4.1.2 Agent Architecture and Platform Governance

The way foundation models are governed, whether proprietary or open-weight, shapes how agentic capabilities spread and who retains control over them. Proprietary models make adoption easier by sparing users the costs and burden of running the underlying infrastructure themselves. That convenience can create a form of dependence. Users may not be able to tailor model weights to private data, may have limited control over updates, and may rely on API access controlled by the provider. In practice, those limits can affect whether downstream agents perform complex tasks consistently. Open-weight models reduce one barrier by making model parameters directly available. But they do not remove the practical difficulties of deployment. Instead, the main constraints shift to compute, engineering capacity, and implementation expertise, which means that well-resourced actors still have a clear advantage. Neither governance model is automatically more equitable. Each comes with its own pattern of access constraints, dependency risks, and quality differences (cf. Tiwana 2014).

4.1.3 Integration and Deployment

Much of an agent’s value depends on what it can connect to in the wider digital environment, including APIs, proprietary databases, and software systems. That gives the actors controlling this “agent infrastructure”, such as cloud providers, financial institutions, and large software platforms, an important gatekeeping role. Through API policies, data-sharing terms, and platform rules, they can widen or narrow the benefits that agents provide to different users and developers. This matters directly for agent *quality*, especially when effective tool use is central to performance (Chan et al., 2025; Rothschild et al., 2025). These chokepoints can create new forms of dependence, although open standards and interoperability requirements may soften them.

4.2. Socio-Political and Institutional Forces

4.2.1 Economic Incentives and Market Dynamics

Market incentives are likely to pull in both directions. Firms have strong commercial reasons to offer basic agentic systems widely and cheaply in order to build large user bases. At the same time, familiar strategies such as product differentiation and price discrimination are likely to

create substantial differences in the *quality* and *quantity* of agents. Premium tiers may provide better planning, broader tool use, or more reliable performance. Pricing may also put large agent swarms, which matter for scalable delegation, beyond the reach of individuals and smaller organisations. The result may be broad formal access alongside continued inequality in the most valuable uses.

4.2.2 Digital Literacy and Human-Agent Interaction

How well people use AI agents will depend in part on existing differences in human capital. In the near term, users may need considerable skill to specify complex goals and manage multi-step delegation, which is likely to favour those with higher levels of digital literacy. Over time, however, more capable agents may reduce the need for specialised technical skill by allowing users to delegate complex tasks through natural language and more intuitive interfaces (Gabriel et al., 2024). Evidence of relatively high usage of LLMs in some middle-income countries, including Nigeria and India (Similarweb, 2026; Semrush, 2026), suggests that language-based interfaces can broaden adoption in ways that more technical interfaces often do not.

4.2.3 Geopolitics and Jurisdictional Fragmentation

At the global level, agentic inequality may be shaped by a tension between techno-national competition and the need for international coordination. National security concerns are already driving protectionist policies, such as semiconductor export controls, which restrict access to foundational AI hardware (Leicht, 2025; Larsen, 2022). As agentic systems mature, similar logics could extend from hardware to software, leading to mercantilist strategies where firms offer degraded export versions of their agents that have weaker planning abilities or tighter limits on autonomous action. This could produce durable international asymmetries in agentic power. Jurisdictional fragmentation may deepen those asymmetries, as divergent regulations on issues such as data privacy create compliance burdens that large, multinational firms are better positioned to absorb than smaller competitors (Mueller, 2021). Against this, an emerging international consensus around AI safety and responsible deployment may provide at least a partial basis for more coordinated and equitable governance (Concordia AI, 2024; Bengio et al., 2025).

4.2.4 Governance, Regulation, and Societal Adaptation

Any response to agentic inequality is constrained by the “pacing problem”: technological change may advance much faster than institutions can adapt (Marchant, 2011). That lag creates a window in which inequalities in the ability to deploy autonomous systems can become entrenched by default, especially where regulatory capture is also a risk (Stigler, 1971). Yet this same transition creates an opening for anticipatory governance. The challenge is to design institutions capable of shaping the deployment of systems that act autonomously in complex social environments before unequal patterns harden into durable structures (Guston, 2014; Floridi et al., 2018).

5. Governing Agentic Inequality: Complexity and Future Directions

5.1. The Governance Challenge

Governing agentic inequality is not only about preventing the harms of disparity; it is also about creating the conditions under which agents can function as a powerful equalising tool. This is a difficult governance problem because it begins with a normative question: what distribution of agentic capability should count as fair? (Gabriel, 2022). Societies already tolerate large disparities in access to expert human services such as legal or financial advice. The question, then, is whether inequalities in agentic capability should be judged by the same standards as disparities in human capability, or whether their scalability and autonomy require stricter standards of justice. A full account of agentic equality lies beyond the scope of this paper. Two directions nevertheless seem especially promising: ensuring universal structural availability as a threshold condition, and designing agent infrastructure with public-interest obligations.

Even if some normative consensus emerges, substantial practical obstacles remain. Existing legal frameworks are poorly suited to harms produced by unequal agentic capability, where disadvantage may arise through cumulative strategic asymmetries rather than a single clearly cognisable wrong (Kolt, 2025). Attempts to level the playing field through direct intervention also face serious coordination problems and the Collingridge dilemma: acting early is difficult because uncertainty is high, while acting late risks the technology already being entrenched

(Collingridge, 1980). Caps on agent power, for example, could constrain innovation, while mandates to redistribute compute could divert resources away from large-scale projects that may offer substantial societal benefits. Government efforts to guarantee universal access through a single provider could produce vendor lock-in and weaken competition, reinforcing rather than reducing inequality (Sastry et al., 2024; Stigler, 1971).

5.2. A Forward-Looking Research Agenda

The path forward requires research organised around each of the paper’s three core dimensions. The goal is to build the knowledge needed to govern disparities around each.

5.2.1 Empirical Foundations: Measuring the Dimensions of Agent Inequality

A first priority is measurement. Before effective interventions can be designed, researchers need robust metrics to track disparities in agent *availability*, *quality*, and *quantity* in real-world settings. Many of the technologies likely to drive this inequality, such as automated negotiation or scaled task delegation, remain novel and largely unmeasured. Key questions include: What are the best methods for empirically tracking the deployment of AI agents and their distributional impacts? How can we conduct direct technical comparisons to assess the extent to which disparities in agent *quality* allow more capable agents to persuade, manipulate, or out-negotiate weaker ones? Possible indicators include: agent deployment rates by income decile and geography, benchmark performance gaps between free and premium agents on standardised task suites, and performance differentials in controlled negotiation settings across capability tiers.

5.2.2 Normative Foundations: Defining Fairness Across the Dimensions

A second programme is normative. Once empirical patterns are clearer, the next task is determining which disparities in agent *availability*, *quality*, and *quantity* are socially or ethically unacceptable. Key questions here include: Which participatory methods, such as citizen assemblies, are best suited to eliciting public preferences about unacceptable capability gaps? How can these preferences be translated into concrete design principles or regulatory guardrails governing disparities in agent *quality*?

5.2.3 Technical and Infrastructural Levers for Equality

A third programme concerns the design of the agentic ecosystem itself. The aim is to build equity-supporting features directly into technical and infrastructural arrangements, reducing harmful capability gaps and limiting excessive concentration in the deployment of large numbers of agents. Key questions include: How to govern critical resources such as compute in ways that balance innovation with equitable access? Which interoperability standards may be needed to prevent platform lock-in from entrenching disparities in agent *quality* (Chan et al., 2025; Kapoor et al., 2025)?

5.2.4 Public Service Models for Agent AI

A fourth line of inquiry concerns non-market models for securing equity. This includes policies intended to guarantee universal access and a sufficient baseline of agent *quality*, while recognising that government-led digital provision has a mixed track record (Dunleavy et al., 2006). One promising idea is “universal basic agency”: the view that access to AI agents should be treated as a primary good in Rawls’s sense, and thus as something all citizens are entitled to as a matter of justice (Rawls, 2001). On this account, agents should function as all-purpose means, enabling people to pursue a wide range of ends and conceptions of the good life. Justice therefore requires more than formal *availability* alone. People must be able to access and use agents of sufficient *quality* for that capability to be genuinely effective across the population, even where others can deploy more powerful or more numerous systems. Key questions include: Should a government provide a public-option AI agent, or would subsidising private services be more effective? What level of agent *quality* would constitute a meaningful baseline of empowerment? How could public agents, or public infrastructure for private agents, be designed to serve public rather than commercial interests? Early proposals for public AI infrastructure suggest that this is a live policy question in several jurisdictions (Jackson et al., 2024; Sitaraman and Parek, 2025).

5.2.5 Regulatory Frameworks for Agent Interactions

A final programme concerns the regulation of asymmetries in agent *quality* and *quantity*. These new asymmetries strain existing legal doctrines,

and without clearer rules, the advantages conferred by superior agents may become entrenched. Key questions include: How can legal frameworks be adapted to provide recourse for individuals harmed by disparities in agent *quality* (Calvo et al., 2020)? What agile or outcome-based regulatory models can keep pace with technical change while protecting people from exploitation by actors deploying vast numbers of agents (Marchant, 2011)?

5.3. Conclusion

Agentic inequality is a new societal challenge produced by the interaction of technical, economic, and political forces. The framework developed here – *availability*, *quality*, and *quantity* – marks out the territory. Scalable delegation represents a genuinely novel source of inequality, as does the shift from contests between people to contests between their agents. However, many of the drivers of inequality are also the potential levers for shaping more equitable outcomes, from model architecture to market design. That window will not remain open indefinitely. Once patterns of agentic deployment become entrenched, they will be far harder to reshape. The research agenda proposed above is an attempt to identify what needs to be understood before that point.

Acknowledgements

We thank Alan Chan, Noam Kolt, Seth Lazar, and Sam Manning for helpful feedback. This work was supported by the Future Impact Group and the Cooperative AI Foundation.

Competing interests

M.S. is employed by Access Partnership, a technology policy consultancy whose clients include companies operating in the AI industry. I.G. is employed by Google DeepMind. O.B. and L.H. declare no competing interests.

References

- Acemoglu, Daron and Restrepo, Pascual. Robots and jobs: Evidence from US labor markets. *Journal of Political Economy*, 128(6):2188–2244, 2020.
- Bengio, Yoshua et al. International AI safety report 2025, 2025. <https://internationalaisafetyreport.org/publication/international-ai-safety-report-2025>.
- Bergman, A. S., Shelby, Renee, and Gabriel, Iason. Access and opportunity. In *The Ethics of Advanced AI Assistants*, pages 146–156. 2024. Preprint at arXiv:2404.16244.
- Besiroglu, Tamay, Bergerson, Sage A., Michael, Amelia, Heim, Lennart, Luo, Xueyun, and Thompson, Neil. The compute divide in machine learning: A threat to academic contribution and scrutiny?, 2024. Preprint at arXiv:2401.02452.
- Bommasani, Rishi, Hudson, Drew A., Adeli, Ehsan, Altman, Russ, Arora, Simran, et al. On the opportunities and risks of foundation models, 2022. Preprint at arXiv:2108.07258.
- Bovens, Mark and Zouridis, Stavros. From street-level to system-level bureaucracies: How information and communication technology is transforming administrative discretion and constitutional control. *Public Administration Review*, 62(2):174–184, 2002.
- Brynjolfsson, Erik and McAfee, Andrew. *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*. W. W. Norton & Company, New York, NY, 2014.
- Brynjolfsson, Erik, Chandar, Bhargav, and Chen, Rui. Canaries in the coal mine? Six facts about the recent employment effects of artificial intelligence. Technical report, Stanford Digital Economy Lab, Stanford Institute for Human-Centered Artificial Intelligence, 2025a.
- Brynjolfsson, Erik, Li, Danielle, and Raymond, Lindsey. Generative AI at work. *The Quarterly Journal of Economics*, 140(2):889–942, 2025b.
- Calvo, Rafael A., Peters, Dorian, Vold, Karina, and Ryan, Richard M. Supporting human autonomy in AI systems: A framework for ethical enquiry. In Burr, Christopher and Floridi, Luciano, editors, *Ethics of Digital Well-Being*, volume 140, pages 31–54. Springer International Publishing, Cham, 2020.
- Chan, Alan, Wei, Kevin, Huang, Sihao, Rajkumar, Nitarshan, Perrier, Elija, Lazar, Seth, Hadfield, Gillian K., and Anderljung, Markus. Infrastructure for AI agents. *Transactions on Machine Learning Research*, 2025. <https://openreview.net/forum?id=Ckh17xN2R2>.
- Collingridge, David. *The Social Control of Technology*. Frances Pinter, 1980.

- Concordia AI. The state of China-Western track 1.5 and 2 dialogues on AI. AI Safety in China, 2024. <https://aisafetychina.substack.com/p/the-state-of-china-western-track>.
- Cui, Kevin Zhiqi et al. The effects of generative AI on high-skilled work: Evidence from three field experiments with software developers. *Management Science*, 2026. <https://doi.org/10.1287/mnsc.2025.00535>.
- Dunleavy, Patrick, Margetts, Helen, Bastow, Simon, and Tinkler, Jane. *Digital Era Governance: IT Corporations, the State, and e-Government*. Oxford University Press, 2006.
- Floridi, Luciano, Cowls, Josh, Beltrametti, Monica, Chatila, Raja, Chazerand, Patrice, Dignum, Virginia, Luetge, Christoph, Madelin, Robert, Pagallo, Ugo, Rossi, Francesca, Schafer, Burkhard, Valcke, Peggy, and Vayena, Effy. AI4People — an ethical framework for a good AI society: Opportunities, risks, principles, and recommendations. *Minds and Machines*, 28(4):689–707, 2018.
- Gabriel, Iason. Toward a theory of justice for artificial intelligence. *Daedalus*, 151(2):218–231, 2022.
- Gabriel, Iason, Manzini, Arianna, Keeling, Geoff, et al. The ethics of advanced AI assistants, 2024. Preprint at arXiv:2404.16244.
- Gimbel, Martha, Kinder, Molly, Kendall, Jake, and Lee, Michael. Evaluating the impact of AI on the labor market: Current state of affairs. Technical report, The Budget Lab at Yale University, 2025.
- Goldberg, Ben et al. AI and the future of digital public squares, 2024. Preprint at arXiv:2412.09988.
- Guston, David H. Understanding anticipatory governance. *Social Studies of Science*, 44(2): 218–242, 2014.
- Hosseini Maasoum, Seyed Mohamad and Lichtinger, Gal. Generative AI as seniority-biased technological change: Evidence from U.S. résumé and job posting data. *SSRN Electronic Journal*, 2025. <https://doi.org/10.2139/ssrn.5425555>.
- Ilves, Luukas, Kilian, Martin, Parazzoli, Stefano Maria, Peixoto, Tiago C., and Velsberg, Ott. The agentic state: Rethinking government for the era of agentic AI. Technical report, Global Government Technology Centre Berlin & The World Bank, 2025.
- Jackson, Ben et al. Public AI: Infrastructure for the common good. Public AI Network, 2024. <https://doi.org/10.5281/zenodo.13914560>.
- Jasanoff, Sheila. *States of Knowledge: The Co-Production of Science and the Social Order*. Routledge, London, 2004.
- Kapoor, Sayash, Kolt, Noam, and Lazar, Seth. Position: Build agent advocates, not platform agents. In *Proceedings of Machine Learning Research*, volume 267, pages 81617–81633, 2025.
- Kasirzadeh, Atoosa and Gabriel, Iason. Characterizing AI agents for alignment and governance, 2025. Preprint at arXiv:2504.21848.
- Kolt, Noam. Governing AI agents, 2025. Preprint at arXiv:2501.07913.
- Korinek, Anton and Vipra, Jai. Concentrating intelligence: Scaling and market structure in artificial intelligence. *Economic Policy*, 40:225–256, 2025.
- Krier, Samuel. Coasean bargaining at scale. AI Policy Perspectives, 2025. <https://www.aipolicyperspectives.com/p/coasean-bargaining-at-scale>.
- Larsen, Bjørn Conrad. Geopolitics of AI and the rise of digital sovereignty. Technical report, Brookings Institution, 2022.
- Lazar, Seth. Governing the algorithmic city. *Philosophy & Public Affairs*, 53(2):102–168, 2025.
- Leicht, Anton. Mercantilist AI policy: Reconstructing the doctrine behind US AI foreign policy. Online essay, 2025. <https://writing.antonleicht.me/p/mercantilist-ai-policy>.
- Li, Junyou, Zhang, Qin, Yu, Yangbin, Fu, Qiang, and Ye, Deheng. More agents is all you need. *Transactions on Machine Learning Research*, 2024.
- Liu, Chang, Wang, Chonghua, Cao, Jiaxuan, et al. A vision for auto research with LLM agents, 2025. Preprint at arXiv:2504.18765.

- Marchant, Gary E. The growing gap between emerging technologies and the law. In Marchant, Gary E., Allenby, Braden R., and Herkert, Joseph R., editors, *The Growing Gap Between Emerging Technologies and Legal-Ethical Oversight*, volume 7, pages 19–33. Springer Netherlands, Dordrecht, 2011.
- Mathur, Arunesh et al. Dark patterns at scale: Findings from a crawl of 11K shopping websites. *Proceedings of the ACM on Human-Computer Interaction*, 3(CSCW):81, 2019.
- Mueller, Benjamin. How much will the artificial intelligence act cost Europe? Technical report, Centre for Data Innovation, 2021.
- Mumford, Lewis. *The Myth of the Machine. Vol. 1: Technics and Human Development*. Harcourt, Brace, Jovanovich, New York, NY, 1967.
- Noy, Shakked and Zhang, Whitney. Experimental evidence on the productivity effects of generative artificial intelligence. *Science*, 381(6654): 187–192, 2023.
- Peng, Sida, Kalliamvakou, Eirini, Cihon, Peter, and Demirer, Mert. The impact of AI on developer productivity: Evidence from GitHub Copilot, 2023. Preprint at arXiv:2302.06590.
- Persily, Nathaniel and Tucker, Joshua A. *Social Media and Democracy: The State of the Field, Prospects for Reform*. Cambridge University Press, 2020.
- Plantin, Jean-Christophe, Lagoze, Carl, Edwards, Paul N., and Sandvig, Christian. Infrastructure studies meet platform studies in the age of Google and Facebook. *New Media & Society*, 20(1):293–310, 2018.
- Prunkl, Carina. Human autonomy at risk? An analysis of the challenges from AI. *Minds and Machines*, 34:26, 2024.
- Rawls, John. *Justice as Fairness: A Restatement*. Harvard University Press, 2001.
- Rothschild, David M., Mobius, Markus, Hofman, Jake M., Dillon, Eleanor W., Goldstein, Daniel G., Immorlica, Nicole, Jaffe, Sonia, Lucier, Brendan, Slivkins, Aleksandrs, and Vogel, Matt. The agentic economy, 2025. Preprint at arXiv:2505.15799.
- Sastry, Girish, Heim, Lennart, Belfield, Haydn, Anderljung, Markus, Brundage, Miles, Hazell, Julian, O’Keefe, Cullen, Hadfield, Gillian K., Ngo, Richard, Pilz, Konstantin, Gor, George, Bluemke, Emma, Shoker, Sarah, Egan, Janet, Trager, Robert F., Avin, Shahar, Weller, Adrian, Bengio, Yoshua, and Coyle, Diane. Computing power and the governance of artificial intelligence, 2024. Preprint at arXiv:2402.08797.
- Selwyn, Neil. Reconsidering political and popular understandings of the digital divide. *New Media & Society*, 6(3):341–362, 2004.
- Semrush. Top websites in Nigeria. <https://www.semrush.com/trending-websites/ng/all>, 2026. Accessed April 2026.
- Sen, Amartya. *Development as Freedom*. Oxford University Press, 1999.
- Similarweb. Top websites ranking for India. <https://www.similarweb.com/top-websites/india/>, 2026. Accessed April 2026.
- Sitaraman, Ganesh and Parek, Ketan. The global rise of public AI. Technical report, Vanderbilt Policy Accelerator, 2025.
- Stigler, George J. The theory of economic regulation. *The Bell Journal of Economics and Management Science*, 2(1):3–21, 1971.
- Summerfield, Christopher et al. The impact of advanced AI systems on democracy. *Nature Human Behaviour*, 9:2420–2430, 2025.
- Tiwana, Amrit. *Platform Ecosystems: Aligning Architecture, Governance, and Strategy*. Elsevier, 2014.
- Tomasev, Nenad et al. Virtual agent economies, 2025. Preprint at arXiv:2509.10147.
- van Dijk, Jan A. G. M. *The Deepening Divide: Inequality in the Information Society*. SAGE, 2005.
- Van Loo, Rory. Digital market perfection. *Michigan Law Review*, 117(5):815–884, 2019.
- Wang, Lei, Ma, Chen, Feng, Xueyang, Zhang, Zeyu, Yang, Hao, Zhang, Jingsen, Chen, Zhiyuan, Tang, Jiakai, Chen, Xu, Lin, Yankai, Zhao, Wayne Xin, Wei, Zhewei, and Wen, Ji-Rong. A survey on large language model

- based autonomous agents. *Frontiers of Computer Science*, 18(6):186345, 2024. URL <http://arxiv.org/abs/2308.11432>.
- Warschauer, Mark. *Technology and Social Inclusion: Rethinking the Digital Divide*. MIT Press, 2003.
- Weidinger, Laura, Mellor, John, Rauh, Maribeth, Griffin, Conor, Uesato, Jonathan, Huang, Po-Sen, Cheng, Myra, Glaese, Mia, Balle, Borja, Kasirzadeh, Atoosa, Kenton, Zac, Brown, Sasha, Hawkins, William, Stepleton, Tom, Biles, Courtney, Birhane, Abeba, Haas, Julia, Rimell, Laura, Hendricks, Lisa Anne, Isaac, William, Legassick, Sean, Irving, Geoffrey, and Gabriel, Iason. Ethical and social risks of harm from language models, 2021. Preprint at arXiv:2112.04359.
- Weyl, E. Glen, Tang, Audrey, and the Plurality Community. *Plurality: The Future of Collaborative Technology and Democracy*. Radicalx-Change, 2024.
- Xu, Jiawei et al. Rethinking the value of multi-agent workflow: A strong single agent baseline, 2026. Preprint at arXiv:2601.12307.
- Yang, Yifan et al. Understanding agent scaling in LLM-based multi-agent systems via diversity, 2026. Preprint at arXiv:2602.03794.
- Zhang, Eric, Zhu, Eric E., Bansal, Gagan, Fournery, Adam, Mozannar, Hussein, and Gerrits, Jacob. Optimizing sequential multi-step tasks with parallel LLM agents, 2025. Preprint at arXiv:2507.08944.
- Zhu, Shijia, Sun, Jingbo, Nian, Yi, South, Tobin, Pentland, Alex, and Pei, Jian. The automated but risky game: Modeling and benchmarking agent-to-agent negotiations and transactions in consumer markets, 2025. Preprint at arXiv:2506.00073.
- Żywot, Aleksander, Chen, Xiang, and de Rijcke, Maarten. Can small agent collaboration beat a single big LLM?, 2026. Preprint at arXiv:2601.11327.