
Position: The Pace of AI Innovation Demands Matching Urgency in Societal Impact Research to Shape Economic Policy

Anonymous Authors¹

Abstract

This position paper argues that society must pursue research into artificial intelligence's (AI) societal impacts with the same urgency that characterizes current technical development. While billions flow into advancing AI capabilities, our understanding of potential societal transformations remains dangerously incomplete. Asymmetric adoption of AI across industries, driven by varying regulatory pressures and institutional constraints, threatens to create economic disruption before society can adapt. As AI capabilities advance, we risk fundamental changes to labor markets and educational systems before developing adequate adaptation strategies. AI's unique Network effects and rapid implementation capabilities create unprecedented challenges for traditional institutional adaptation. Drawing detailed parallels between Major League Baseball (MLB's) economic structure and emerging patterns in AI-automated knowledge work, we identify mechanisms that could accelerate opportunity concentration, while exploring educational institutions' vulnerability to these changes and identifying critical research priorities.

1. Introduction

The artificial intelligence (AI) narrative has shifted dramatically from augmenting human capabilities to positioning autonomous agents as direct replacements for knowledge workers. **The rapid pace of AI innovation demands equally urgent research into societal impacts to inform policy interventions that preserve economic mobility and protect individual opportunity.**

Modern AI models have demonstrated consistently improved performance on professional benchmarks. For example, OpenAI's GPT-4 scored in the top 10th percentile on the US Bar Exam, a significant improvement from GPT-3.5's bottom 10th percentile performance (Achiam et al., 2023). Similarly, Google's Med-Gemini model achieved 91.1% accuracy on the MedQA dataset in May 2024, an

improvement from 67.2% in December 2022 (Saab et al., 2024). While the presence of these assessments in training data complicates interpretation, the trend toward increasing autonomous capabilities in knowledge work sectors is clear.

Major technology companies have invested billions into AI research and development. In the last quarter of 2024, Anthropic, xAI and OpenAI raised billions of dollars (Amazon, 2024; Achiam et al., 2023; Wiggers, 2024). Venture capital has followed suit, as AI startups now command 46.4% of the \$209 billion in total venture funding, up from under 10% a decade ago (Hu, 2025). While these massive investments are framed around long-term aspirations of artificial general intelligence (AGI), the near-term path to financial returns likely involves replacing expensive human knowledge workers with less costly autonomous systems.

Consequently, if autonomous agents can perform existing tasks more cheaply, basic market forces will drive human labor reduction. Corporations facing pressure to justify massive AI investments are unlikely to maintain redundant human workforces. This creates a self-reinforcing cycle, as companies increasingly view AI systems as a replacement rather than a tool, they restructure business models and investment strategies accordingly, further accelerating technical capabilities and reinforcing the replacement narrative.

Discussions of AI's impact often polarize between dismissing it as hype and envisioning a future of abundance. While the narrative around AI safety centers around alignment and catastrophic risks, the gap between technological innovation and societal adaptation (Kane, 2017) is often overlooked. This paper examines the crucial period where the uneven adoption of increasingly capable AI systems across industries, driven by varying market pressures and regulations, threatens to reshape knowledge work before institutions and policies can adapt.

To understand the potential future of knowledge work under AI automation, we can look to an established model of extreme market concentration: Major League Baseball (MLB). The economics of professional baseball, with its stark division between major and minor leagues, provides a mature example of how winner-take-most markets develop

and stabilize (McMahon Jr & Abreu, 1998). This parallel provides a framework for understanding how AI could reshape professional hierarchies and economic mobility. The parallel between MLB's economic structure and potential AI-driven knowledge work transformation has inherent limitations, as MLB represents an extreme version of this type of market economy. Knowledge work encompasses diverse domains with complex value creation mechanisms that will react to automation differently. However, examining MLB's mature winner-take-most market structure provides valuable insights into how and why similar economic patterns could emerge in knowledge work, and how even modest shifts in economic concentration could have broad societal impact. This paper is intended to be a call to action for policymakers, researchers in AI and related social sciences, and industry leaders who must grapple with the potential consequences of these technological shifts.

This transformation threatens not just current knowledge workers but the entire educational and professional infrastructure built around knowledge work careers. The following sections examine how this cycle could manifest through direct economic incentives, network effects, and fundamental changes in professional opportunity structures, and draws a parallel to how knowledge working professions could adopt MLB's winner-take-most economic model that increasingly reflects industry level trends at the firm level (Autor et al., 2020).

2. Professional Sports Economics as a Model for Knowledge Work Transformation

The MLB model exemplifies a mature winner-take-most market structure with extreme economic stratification (McMahon Jr & Abreu, 1998). In 2024, Major League players earned average base salaries of almost \$5 million annually, while Triple-A players, one step from the Major Leagues, started at just \$35,000 (Statista, 2024; Cooper, 2024b). This dramatic compensation cliff, where being the 700th versus 1000th best player means earning 15-20 times more, illustrates how winner-take-most markets develop and stabilize. The path includes intense selectivity as only 0.5% of high school players reach professional baseball, and just 17.6% of drafted players make the Major Leagues (Web, 2024; Cooper, 2024a), with research suggesting that the economic viability of an MLB career is limited to those drafted in the top 171 picks (Pifer et al., 2020).

2.1. AI's Amplification of Winner-Take-Most Dynamics

While MLB's structure reflects supply-demand dynamics in sports, AI automation could create more pronounced effects in knowledge work. Unlike MLB's average career tenure of 5.6 years (Witnauer et al., 2007), knowledge workers can maintain productivity longer, allowing incumbents to accu-

mulate and retain advantages. Where baseball's hierarchy is performance-based, AI could create more arbitrary divisions through early adoption, institutional positions, and networks. These less predictable factors could create more extreme opportunity concentration. The next section examines how market forces could drive knowledge work toward MLB's winner-take-most structure, with more severe implications for economic mobility.

3. The AI Automation of Knowledge Work: Market Dynamics and Technological Convergence

AI automation transforms knowledge work through three interconnected mechanisms: economic incentives driving adoption, network effects accelerating advantage concentration, and resulting changes in professional opportunity distribution. Understanding their interaction is crucial for anticipating the pace and scope of change.

3.1. Corporate Incentives Towards Automation

While automation pressure varies across knowledge work domains based on task complexity and expertise, these roles are susceptible to AI disruption due to their reliance on information processing and decision-making systematized through natural language interfaces. Agentic AI, defined by Bousetouane, refers to "autonomous, intelligent systems powered by LLMs that integrate modular components—reasoning, memory, cognitive skills, and tools—to solve complex tasks in dynamic and evolving environments" (Bousetouane, 2025). This is a broad definition, and specific tasks will inherently have varied levels of human involvement and orchestration even with increasing automation, but conceptually this definition aims to capture shifting work from people to AI-based automation.

Recent market developments highlight this transition towards full automation, including increases in corporate investment and shifting marketing strategies. "Vertical AI Agents" have become a common description of autonomous entities executing entire workflows or tasks through natural language requests tailored for specific industries (Bousetouane, 2025). This differs from software-as-a-service, which managed repetitive business tasks with clear boundaries using structured data (Bousetouane, 2025).

Major tech companies are rapidly positioning AI as a workforce replacement. Salesforce's "Agentforce" promises "a limitless AI workforce," with CEO Marc Benioff citing 30% engineering productivity gains eliminating 2025 engineering hiring needs (Martin, 2024; Salesforce, 2024). Cognition AI launched "Devin", an autonomous software engineer (Cognition Labs, 2024), while Mark Zuckerberg projects mid-level engineer-equivalent AI by 2025 (Varanasi, 2025).

Artisan, an AI-powered sales automation platform, launched a provocative marketing campaign in 2024 urging companies to "Stop Hiring Humans." The company later acknowledged the intentionally controversial nature of this messaging (Carmichael-Jack, 2024), but the campaign's resonance reflects a shift in how AI companies position their products. In a recent podcast, Y Combinator partners projected that the market potential for vertical AI agents could exceed traditional Software-as-a-Service (SaaS) models by an order of magnitude, specifically because these AI-driven solutions promise to reduce or eliminate human operational overhead in ways that SaaS products could not (Louise, 2024). In an interview in January 2025, Anthropic CEO Dario Amodei projected that AI will be better than humans at virtually all tasks by 2027, or shortly thereafter (Edwards, 2025). A Bloomberg Intelligence study estimated that automation driven by AI could directly lead to 200,000 job cuts in the banking industry in the next three to five years, illustrating the scale of potential workforce transformation (Kelly, 2025).

This shift creates concerning implications for knowledge worker wages and employment. Research examining AI's impact on freelance workers across online labor markets revealed that workers initially benefited from AI capability increases through enhanced productivity, but after reaching a capability "inflection point," workers became worse off economically with further AI improvements (Qiao et al., 2023). Korinek and Suh's research presents a framework for understanding potential wage trajectories. If human work has unlimited complexity, wages could theoretically rise indefinitely as AI advances. However, if human cognitive capabilities are bounded, wages would initially surge as machines complement human labor, but eventually collapse as AI systems become capable of performing increasingly complex cognitive tasks independently (Korinek & Suh, 2024).

Successfully automating knowledge work roles presents companies with significant profit opportunities. This automation pressure may become self-reinforcing across industries. Even organizations initially committed to maintaining human workforces may face competitive pressure to adopt automation when rivals demonstrate significant cost advantages (Vale, 2023). This dynamic parallels how previous efficiency innovations, like outsourcing, spread across industries not always because companies wanted to eliminate domestic jobs, but because market competition left them little choice (Kulembayeva et al., 2022). While AI-based automation remains nascent and unproven, the general narrative surrounding efficiency gains notably lacks serious examination of broader societal implications for knowledge work and human labor in general if these strategies prove successful at scale.

3.2. A Potential Knowledge Work Stratification

The economic stratification visible in MLB reflects a mature model of how winner-take-most markets develop and stabilize over time (Rosen, 1981). Just as MLB has evolved clear tiers of professional opportunity, the AI-driven transformation of knowledge work could create similar hierarchical structures, but with even more profound implications given its broader economic scope.

The resulting knowledge work economic structure could mirror baseball's three-tier system across various industries. At the top, "Major League" knowledge workers would command premium compensation for handling novel challenges and high-stakes decisions, or having economic control over many influential automated agents. An example would be a diagnostic medical doctor that specialized in rare diseases that are not typically captured data distribution of LLMs. Another example would be developers or corporations that are able to automate functions previously performed by humans but at a significantly greater scale, such as a single corporation that can replicate the output of entry-level lawyers.

A middle tier of "Minor League" positions might exist for managing AI systems and handling edge cases, but with significantly lower compensation than the top tier. This could include auditing or supervising agentic workflows for specific applications, such as a paralegal in charge of ensuring all case research is factual. A key distinction between these positions versus the "Major League" positions is that and the skills required to execute these tasks will not be highly specialized, and may be subject to intense competition.

It is unclear the salary level this type of position may be able to demand, but if there is intense competition due to the large supply of workers relative to available jobs, it is possible that this will depress salaries for this category of work, mirroring minor league salaries observed in MLB. Technological progress and individual industry regulation and adoption will draw the distinction between how prevalent "Major" and "Minor" league positions that will exist in specific knowledge working fields, but conceptually the more capable and accepted autonomous agents become, the more leverage those at the top will have to put pressure on knowledge worker employment and salaries in this category of work.

Existing professions will see reduced available roles due to the automation of previously human tasks, including scenarios with significantly reduced human oversight, such as one manager overseeing six autonomous software agents instead of six human developers. Knowledge workers who cannot secure positions in the top tiers may need to transition to entirely different fields, similar to how most American baseball players must pursue alternative careers when their

collegiate or amateur career concludes.

While competitive effects already exist in knowledge work, with varying compensation reflecting market dynamics and skill differentials, AI automation could dramatically amplify these disparities. Our current societal structure positions these knowledge working professions as safe and desirable, and if that changes rapidly, a large swath of the population will be left unprepared to replace their desired or current career with a viable alternative. While the MLB model developed in a contained ecosystem affecting a relatively small population, there are an estimated one billion knowledge workers globally and over 100 million in the United States as of 2021 (Luna-Ostaseski, 2021). The AI-driven transformation of knowledge work could affect hundreds of millions of professionals and reshape entire industries, potentially undermining one of society's primary mechanisms for economic mobility and stable employment.

3.3. Network Effects and Innovation Speed

AI's unique network effects accelerate knowledge work stratification beyond traditional career hierarchies. While advancement typically depends on individual skill development, AI creates compounding advantages through data accumulation and system optimization (Levine & Jain, 2023). In fields like legal services, early adopters build insurmountable leads as standardized work products generate valuable training data (Hagiu & Wright, 2020). Large foundation model providers further leverage their data and algorithmic expertise across industries, creating self-reinforcing advantages through user interactions (Iansiti & Lakhani, 2020). User queries generate valuable training data, creating compounding advantages for early adopters through system improvements (Constantinides et al., 2018).

The speed of potential AI implementation represents another crucial dynamic. Traditional institutional structures, particularly in government agencies and heavily regulated industries, often require extensive procurement and approval processes for new technology adoption. These processes, designed for an era of slower technological change, may create insurmountable disadvantages as more agile organizations rapidly iterate on their AI implementations (Clark & Hadfield, 2019). This is particularly true in AI, where government entities may lack internal expertise and the solidified standards and guidance needed to manage the procurement process independently and efficiently (Johnson et al., 2024; Medaglia et al., 2023; Zick et al., 2024). A startup capable of near-constant AI system refinements could accumulate months or years of practical advantages over organizations bound by extensive regulation.

These advantages transcend technological capability, as knowledge workers in AI-forward environments gain disproportionate opportunities to develop expertise and pio-

neer novel applications. This human capital development, combined with improving AI systems, creates another feedback loop where the most capable practitioners gravitate toward and further improve the most advanced implementations. The resulting concentration of both human and artificial capability may accelerate the transition toward the winner-take-most economic dynamics observed in professional sports, but with even stronger barriers to new entrants.

Network effects create rigid divisions between knowledge work tiers. "Major League" workers controlling AI systems extend their influence through compounding data advantages. The "Minor League" tier faces increasingly critical but constrained oversight roles with limited advancement opportunities. Those in automated domains encounter growing barriers to upward mobility as technical and organizational gaps widen.

4. Workforce Impact Analysis: The Transformation of Knowledge Work Economics

4.1. Quantitative Impact Projections - An Example

The legal profession illustrates the potential scale of this transformation. Currently, the United States has approximately 1.3 million practicing lawyers, with law schools producing roughly 35,000 new graduates in 2023, and a 5% unemployment rate for those graduates (Association, 2024a;b). Market analysis from Clio suggests that up to 74% of hourly billable tasks could be automated with AI, and AI usage has jumped to 79% of legal professionals in 2024 compared to 19% in 2023. A Goldman Sachs study estimated that AI could automate 44% of legal tasks and almost 40% of legal market jobs (Clio, 2024; Hatzius et al., 2023).

Comparable or even diminished dynamic shifts in the legal field toward the MLB economic model could impact hundreds of thousands of professionals. Based on these projections, a rough simulation suggests that a 33% industry contraction would leave 870,000 available legal positions. This contraction directly threatens entry-level legal roles, as modern LLMs excel in tasks like documentation and data analysis traditionally performed by junior associates (Clio, 2024). This automation pressure could lead to reduced hiring, downward pressure on entry-level salaries, and intensified competition for a shrinking number of entry-level positions. This impact would likely compound annually, as it is logical to expect educational patterns to lag hiring effects, such that there is a glut of unemployed graduates that builds up over time until the market corrects itself. Such a reduction would fundamentally alter the economics of legal education and career planning, creating intense competition for remaining "Minor League" positions and likely apply

downward pressure on salaries. Given the significant cost associated with legal education (Wang, 2024), this could have significant economic consequences for individuals who are unable to obtain income that justifies their initial educational investment.

While automation potential varies across professions, the legal sector illustrates potential workforce transformation scale. This pattern could replicate across industries, with displaced workers competing for fewer opportunities. Some examples include software engineering shifting to high-level architecture, financial services to complex strategy, and medicine to difficult cases and human-centric care, while AI handles routine tasks.

While it is unclear exactly how much automation will apply to each individual industry, and the MLB represents an extreme example of a winner-take-most market, a shift in this direction and downward pressure on job and salary opportunities in knowledge work professions would have a clear economic impact on individuals and have significant implications for upward societal mobility.

4.2. The Unique Nature of AI-Driven Technological Change

Previous technologically driven revolutions created new work categories (Hötte et al., 2023). The first industrial revolution increased productivity through mechanical production (Groumpos, 2021), while the second created engineering and management roles. The digital revolution systematized procedures while generating demand for knowledge workers. Each of these revolutions enhanced productivity and directly led to increases in the standard of living, new careers related to that new technology, and in the case of the first industrial revolution led to sustained population growth for the first time in history (Groumpos, 2021).

Despite these positive outcomes viewed from a historic perspective, there was significant push-back at the time, most famously exemplified by the Luddite movement's destruction of machinery during the first industrial revolution (Groumpos, 2021; Hötte et al., 2023; Acemoglu & Restrepo, 2019). Acemoglu and Restrepo assert that in the past, some technologies have displaced labor due to automation, while others led to new tasks and careers. Importantly, they emphasize that the "presumption that *all* technologies increase (aggregate) labor demand simply because they raise productivity is wrong" (Acemoglu & Restrepo, 2019).

While AI adoption will inevitably create new roles such as AI agent supervisors, its fundamental characteristics suggest a departure from historical patterns of technological disruption. AI presents three distinctive dynamics that challenge traditional assumptions about technological progress, departing from historical patterns where automation created

new forms of human work.

First, AI directly replicates human cognitive work rather than simply automating physical or routine tasks. AI systems are routinely compared directly to human equivalents, and are increasingly being designed to completely replace knowledge workers by performing the same function in a more cost-effective manner. Second, the unprecedented pace of AI innovation suggests workforce displacement could occur in years rather than decades, potentially outpacing sovereign, regulatory and education entities ability to react to the changes. Third, AI's unique capacity for recursive self-improvement through automating cognitive development itself creates the potential for compounding capabilities that accelerate beyond human's ability to adapt. For instance, AI systems can now assist in their own development through code generation and optimization, creating a feedback loop of improvement that has no clear parallel in previous technological revolutions.

This shift transcends traditional automation of lower cognitive tasks, potentially eliminating human involvement across entire categories of knowledge work. This threatens existing paradigms like Skill-Based Technical Change (SBTC), which suggests that that education and skill development provided reliable paths to economic security, as new technologies increase demand for skilled workers while decreasing demand for unskilled labor, like how computers increased demand for knowledge workers that could use them effectively (Violante, 2008).

AI's ability to perform complex cognitive tasks threatens the core assumption of SBTC. At minimum, it risks concentrating SBTC's traditional benefits among an increasingly small number of beneficiaries, similar to the MLB economic model. Unlike previous technological transitions where human capital could be reliably redirected toward higher-value activities, AI's broad cognitive capabilities may leave fewer clear paths or opportunities for human specialization, as part of the core functionality of modern LLMs is to reflect the median level of human expertise across a wide variety of tasks.

While some argue that AI could follow historical patterns, driving productivity gains and creating as many or more jobs than it displaces, the unprecedented combination of cognitive replication, rapid advancement, and recursive improvement suggests this transition may follow fundamentally different patterns. Aligning with this paper's central thesis, society must prepare for scenarios where traditional knowledge work opportunities contract significantly. This dynamic is particularly punishing given knowledge work professionals typically invest heavily in education and specialized training under the assumption that their skills would retain or increase in value over time. In such a future, those already holding positions of authority, such as partners at

law firms, executives, or successful entrepreneurs, could leverage AI to expand their influence while reducing dependence on junior talent, accelerating economic concentration among a smaller elite (Korinek, 2024).

5. Alternative Views

5.1. Technological Optimism and the Creation of New Opportunities

A prevalent viewpoint, rooted in historical precedent, posits that technological advancements, while initially disruptive, ultimately lead to net job creation. This perspective draws parallels to previous industrial revolutions, where automation of existing tasks led to the emergence of entirely new industries and professions (Acemoglu & Restrepo, 2019). For example, the advent of computers created demand for software developers, data analysts, and IT professionals, roles that were previously unimaginable. Proponents of this view argue that AI will follow a similar trajectory, generating demand for AI trainers, data labelers, AI ethicists, prompt engineers, AI maintenance technicians, and other yet-to-be-defined roles.

Furthermore, this perspective emphasizes the potential for AI to augment human capabilities rather than replace them entirely. By automating routine and repetitive tasks, AI could free up human workers to focus on more creative, strategic, and interpersonal aspects of their work (Deming, 2017; Bessen, 2019).

5.2. Market Self-Regulation and the Limits of Intervention

Another counterargument suggests that market forces will naturally adapt to AI advancements, minimizing the need for significant policy intervention. Urgent societal research might be unnecessary or even counterproductive if it leads to regulatory constraints that slow beneficial innovation, or if the public sector does not have resources to keep up with the pace of private-sector driven innovation (Clark & Hadfield, 2019). It is important to strike a balance between innovation and safety, and applying too much policy or regulatory pressure could stifle innovation relative to other sovereign competitors (Panait et al., 2021). AI could drive broad macro-level productivity gains, and asymmetric adoption across industries and regions provides natural experiments for policymakers.

These optimistic views, while acknowledging the potential for short-term disruption, often underestimate the unique characteristics of AI-driven automation relative to prior technological shifts. AI directly replicates and potentially surpasses human cognitive abilities across a wide range of knowledge work domains, and does so at an accelerating pace. This combination of breadth and speed creates a qual-

itatively different challenge compared to past technological transformations. Furthermore, the argument for market self-regulation overlooks the potential for significant and rapid increases in inequality as market forces alone may not ensure equitable distribution of AI's benefits or adequately address the needs of displaced workers. This potential for rapid and significant disruption to knowledge worker employment may necessitate policy intervention that market forces alone cannot address.

6. Societal Implications and Future Research Needs: Addressing Systemic Risk

6.1. Education Institution Inertia and Market Misalignment

The current educational system's structure creates particular vulnerabilities in this transition. Individual institutions that unilaterally reduce enrollment or radically reshape programs risk losing market position and revenue, while those maintaining traditional approaches despite declining career prospects may contribute to systemic oversupply of graduates. This collective action problem impedes educational system adaptation, even when administrators recognize the need for change. According to the US Census Bureau, the number of adults 25 and older with a Bachelor's degree increased from an estimated 62 million in 2010 to just under 91 million in 2023 (U.S. Census Bureau, 2023).

The current educational financing system amplifies this misalignment by concentrating financial risk on students through loan obligations. Educational institutions face minimal immediate consequences from graduate unemployment, with risks limited to reputation damage and reduced future donations and applications. Student loan debt has ballooned from \$461 billion to \$1.77 trillion from 2006 to 2024, a 284% increase, compared to a 59% increase in CPI inflation over the same time period (U.S. Bureau of Labor Statistics, 2024; Federal Reserve Bank of St. Louis, 2024). This structure could become increasingly problematic as automation reduces traditional career opportunities, and leaves a huge oversupply in professions that previously reflected a seemingly safe return on investment for students.

Current students and recent graduates face urgent challenges as traditional educational and career guidance becomes obsolete, yet institutions lack mechanisms for rapid curricular adaptation. This creates an immediate need for alternative career development frameworks that acknowledge the reality of AI automation. As AI capabilities accelerate, the lag in educational response means that institutions continue producing graduates for rapidly contracting traditional roles. This temporal mismatch threatens to systematically amplify existing socioeconomic disparities by concentrating adaptation capabilities among those with resources to

independently navigate the transition.

The contraction of knowledge work as a reliable path to economic mobility threatens both individual opportunity and broader social stability. Higher education's role as a gateway to knowledge work careers has historically provided a reliable path to economic advancement (Woodhall, 1987; Kim et al., 2018; Violante, 2008). This disruption threatens to unravel decades of social infrastructure built around knowledge work careers. Under AI automation, an expensive knowledge work degree might become a source of insurmountable debt without corresponding opportunity. The inter-generational implications extend beyond individual career trajectories to threaten entire family mobility strategies. Families who structured their children's education around traditionally stable professions like medicine, law, or engineering may find these carefully planned pathways invalidated by automation, with severe financial consequences.

6.2. Research and Policy Imperatives

The pace and scale of AI-driven changes in knowledge work demand immediate research attention across multiple domains. Critically, this challenge extends beyond the scope of AI research alone, it represents a complex socio-technical problem requiring unprecedented collaboration across disciplines. Economists, educators, policy makers, labor experts, and technologists must work together to ensure AI's productivity potential translates into broad societal benefits rather than exacerbating existing inequalities.

Workforce transition analysis must extend beyond traditional skill mapping to understand the persistence and adaptability of cognitive capabilities under increasing automation. While previous technological transitions primarily affected specific task categories, allowing workers to redirect existing skills to adjacent roles, AI automation will affect multiple cognitive domains at varying speeds depending on regulation and corporate appetite for adoption. As an example, modern LLMs may evolve such that they are demonstrably better by any measurable category for medical diagnosis, but given regulatory concerns and general public trust it may take years for widespread adoption of AI to take hold. Contrast that with an industry like software engineering which has limited regulation and has a culture of automation, it is highly likely that the same scenario would lead to close to immediate adoption and disruption (Korinek, 2024).

Economic impact modeling demands refactoring, as rapid, simultaneous transformation across knowledge work sectors creates feedback loops beyond existing models' capacity. Research must examine not only direct employment effects but also the cascade of impacts through supporting industries and services. Geographic distributions of impact may prove particularly important, as AI automation could ac-

celerate existing trends toward concentration of economic opportunity in specific regions or organizations (Korinek, 2024).

Policy framework development presents the most urgent research need, particularly in understanding how to prevent knowledge work from fully replicating professional sports' extreme stratification while preserving innovation incentives. While MLB's market concentration affects a limited population, knowledge work transformation could impact hundreds of millions of careers rapidly. Current policy approaches, designed for gradual technological transitions, prove inadequate for addressing rapid, systemic changes in knowledge work economics (Clark & Hadfield, 2019). The temporal mismatch between AI advancement and traditional policy development cycles creates particular challenges. While AI capabilities can scale globally within months, policy responses typically require years of development and implementation. This mismatch demands new approaches to technology policy that emphasize anticipation and proactive intervention over reactive adjustment. In this US this is particularly true at the Federal level, where inaction may lead to state-level AI legislation becoming default country-level guidance.

Policy interventions, from tax incentives for human-AI collaboration to benchmarks that emphasize complementary rather than competitive performance, could help shape more sustainable labor market outcomes (Bell & Korinek, 2023). Research must focus on identifying which aspects of knowledge work remain persistently valuable as routine cognitive tasks become increasingly automated.

These research imperatives connect directly to broader societal stability concerns. The potential erosion of knowledge work as a reliable path to economic mobility threatens not just individual opportunity but fundamental social narratives about education, merit, and advancement. Research must examine how significant concentration of opportunity might affect social cohesion and democratic stability, while developing policy frameworks that balance innovation and economic efficiency against broader societal welfare considerations.

The machine learning community has a unique responsibility and opportunity to shape this transition. Moving forward, technical AI research cannot exist in a vacuum, it must also explicitly consider potential societal and economic impacts. This includes developing frameworks for human-AI collaboration that preserve economic opportunity, creating benchmarks that measure complementary rather than competitive performance, and investigating approaches that distribute rather than concentrate AI's economic benefits. These directions would help align technical innovation with broader societal welfare while advancing the field's fundamental capabilities.

7. Conclusion

The transformation of knowledge work through AI automation could threaten economic and educational systems that have historically enabled broad-based social mobility. Unlike previous technological revolutions that redistributed labor across sectors, AI automation threatens to disrupt multiple knowledge work domains while consolidating opportunities among a small elite. This transformation, driven by the convergence of technological capability and market incentives, could rapidly reshape professional hierarchies that took generations to establish and cause widespread automation of jobs across variable time horizons depending on industry, which could cause punishing unemployment and societal unrest.

Current educational and professional development systems, optimized for stable career trajectories and requiring significant individual monetary investment, are particularly vulnerable to rapid automation. High educational debt becomes riskier as career paths become less predictable, rigid credentialing requirements clash with the need for rapid adaptation, and traditional professional development models could struggle to keep pace with accelerating technological displacement. Addressing these challenges requires fundamental reforms across educational financing, professional training, and labor market policies to balance innovation incentives with societal stability.

Though the transition timeline remains uncertain, and organizational change typically lags technological advancement (Kane, 2017), major voices inside frontier labs have been insisting that the rate of progress is set to accelerate. Ethan Mollick articulates this mismatch between AI acceleration and the lack of corresponding societal preparedness: "What concerns me most isn't whether the labs are right about this timeline - it's that we're not adequately preparing for what even current levels of AI can do, let alone the chance that they might be correct. While AI researchers are focused on alignment, ensuring AI systems act ethically and responsibly, far fewer voices are trying to envision and articulate what a world awash in artificial intelligence might actually look like (Mollick, 2025)."

Crucially, this transformation reflects deliberate market strategies and incentives rather than pure technological inevitability. Proactive intervention could shape more equitable outcomes, but the window for effective action may be narrowing as market forces accelerate adoption. Rather than treating AI advancement and societal preparation as separate challenges, we must recognize them as fundamentally interconnected. Realizing AI's potential benefits requires parallel investment in human capital development and robust social support systems (Korinek, 2024).

The American social safety net, built around stable

employer-provided benefits, appears particularly vulnerable to these changes. Unlike previous transitions where displaced workers could find similar roles in adjacent industries, AI's broad applicability across knowledge work categories may require fundamentally rethinking how society provides essential benefits and maintains economic security. This challenge extends beyond individual sectors to question basic assumptions about professional education, career development, and economic mobility. The lack of social safety net not only would have a negative impact on individuals who lose jobs to automation, but also is likely to motivate inefficient human labor for self-preservation purposes, rather than leveraging AI for more efficient broad economic gain (Bell & Korinek, 2023).

The MLB comparison in this paper serves as an analytical framework rather than a prediction, illuminating how AI adoption could reshape knowledge work economics across industries. This parallel helps us understand potential patterns of opportunity concentration and economic stratification while acknowledging that proactive policy intervention could shape different outcomes.

This analysis points toward several crucial socio-technical research directions: understanding sector-specific automation trajectories, developing adaptive educational models, and designing policy frameworks that encourage innovation but maintain broad-based economic opportunity in an increasingly automated knowledge economy. Even if estimates for the arrival of AGI are overstated or too early, the consistency of opinion across AI industry experts of increasingly capable AI systems that challenge human intelligence demands that AI research deeply integrate societal impact moving forward. Proactive policy intervention and institutional reform are essential to ensure AI advancement generates broad societal benefits rather than concentrating opportunities within an increasingly narrow population.

References

- Acemoglu, D. and Restrepo, P. Automation and new tasks: How technology displaces and reinstates labor. *Journal of economic perspectives*, 33(2):3–30, 2019.
- Achiam, J., Adler, S., Agarwal, S., Ahmad, L., Akkaya, I., Aleman, F. L., Almeida, D., Altenschmidt, J., Altman, S., and et al., S. A. Gpt-4 technical report, 2023.
- Amazon. Amazon invests additional \$4 billion in anthropic ai, 2024. URL <https://www.aboutamazon.com/news/aws/amazon-invests-additional-4-billion-anthropic-ai>.
- Association, T. A. B. Profile of the legal profession 2024 - demographics, 2024a. URL <https://www.americanbar.org/>

- news/profile-legal-profession/
demographics/.
- Association, T. A. B. Employment outcomes as of march 15, 2024 (class of 2023 graduates), 2024b. URL https://www.americanbar.org/content/dam/aba/administrative/legal_education_and_admissions_to_the_bar/statistics/2024/class-of-2023-online-table.pdf/.
- Autor, D., Dorn, D., Katz, L. F., Patterson, C., and Van Reenen, J. The fall of the labor share and the rise of superstar firms. *The Quarterly Journal of Economics*, 135(2):645–709, 2020.
- Bell, S. A. and Korinek, A. Ai’s economic peril. *Journal of Democracy*, 34(4):151–161, 2023.
- Bessen, J. Automation and jobs: When technology boosts employment. *Economic Policy*, 34(100):589–626, 2019.
- Bousetouane, F. Agentic systems: A guide to transforming industries with vertical ai agents. *arXiv preprint arXiv:2501.00881*, 2025.
- Carmichael-Jack, J. The story behind the ”stop hiring humans” billboards in san francisco, 2024. URL <https://www.artisan.co/blog/stop-hiring-humans>.
- Clark, J. and Hadfield, G. K. Regulatory markets for ai safety. *arXiv preprint arXiv:2001.00078*, 2019.
- Clio. Latest Legal Trends Report shows how AI will transform legal services, 2024. URL <https://www.clio.com/about/press/cliio-latest-legal-trends-report/>.
- Cognition Labs. Introducing Devin, an AI software engineer, 2024. URL <https://www.cognition.ai/blog/introducing-devin>.
- Constantinides, P., Henfridsson, O., and Parker, G. G. Introduction—platforms and infrastructures in the digital age, 2018.
- Cooper, J. How many mlb draftees make it to the majors, 2024a. URL <https://www.baseballamerica.com/stories/how-many-mlb-draftees-make-it-to-the-majors/>.
- Cooper, J. How much are minor league baseball players paid in 2024?, 2024b. URL <https://www.baseballamerica.com/stories/how-much-are-minor-league-baseball-players-paid-in-2024/>.
- Deming, D. J. The growing importance of social skills in the labor market. *The quarterly journal of economics*, 132(4):1593–1640, 2017.
- Edwards, B. Anthropic chief says ai could surpass “almost all humans at almost everything” shortly after 2027, 2025. URL <https://arstechnica.com/ai/2025/01/anthropic-chief-says-ai-could-surpass-almost-all-humans-at-almost-everything/>.
- Federal Reserve Bank of St. Louis. Student loans owned and securitized, outstanding (SLOAS), 2024. URL <https://fred.stlouisfed.org/series/SLOAS>. Retrieved from FRED, Federal Reserve Bank of St. Louis.
- Groumpos, P. P. A critical historical and scientific overview of all industrial revolutions. *IFAC-PapersOnLine*, 54(13): 464–471, 2021.
- Hagiu, A. and Wright, J. When data creates competitive advantage. *Harvard business review*, 98(1):94–101, 2020.
- Hatzius, J. et al. The potentially large effects of artificial intelligence on economic growth briggs kodnani. *Goldman Sachs*, 1, 2023.
- Hötte, K., Somers, M., and Theodorakopoulos, A. Technology and jobs: A systematic literature review. *Technological Forecasting and Social Change*, 194:122750, 2023.
- Hu, K. Ai startups drive vc funding resurgence, capturing record us investment in 2024, 2025. URL <https://www.reuters.com/technology/artificial-intelligence/ai-startups-drive-vc-funding-resurgence-capturing-record-us-investment-in-2024-2025-01-09/>.
- Iansiti, M. and Lakhani, K. R. *Competing in the age of AI: Strategy and leadership when algorithms and networks run the world*. Harvard Business Press, 2020.
- Johnson, N., Silva, E., Leon, H., Eslami, M., Schwanke, B., Dotan, R., and Heidari, H. Public procurement for responsible ai? understanding us cities’ practices, challenges, and needs. *arXiv preprint arXiv:2411.04994*, 2024.
- Kane, G. C. Digital disruption is a people problem, 2017. URL <https://sloanreview.mit.edu/article/digital-disruption-is-a-people-problem/>.
- Kelly, J. 200,000 wall street jobs may be slashed by artificial intelligence, 2025. URL <https://www.forbes.com/sites/jackkelly/2025/01/09/200000-wall-street-jobs-may-be-slashed-by-ai/>.
- Kim, C., Tamborini, C. R., and Sakamoto, A. The sources of life chances: does education, class category, occupation, or short-term earnings predict 20-year long-term earnings? *Society for Sociological Science*, 2018.

- Korinek, A. Economic policy challenges for the age of ai. Technical report, National Bureau of Economic Research, 2024.
- Korinek, A. and Suh, D. Scenarios for the transition to agi. Technical report, National Bureau of Economic Research, 2024.
- Kulembayeva, F., Seitkazyeva, A., and Yelshibayev, R. Economic efficiency of outsourcing business models: A comparative assessment. *Global Journal of Flexible Systems Management*, 23(1):75–88, 2022.
- Levine, S. and Jain, D. How network effects make ai smarter. *Harvard Business Review*, 14, 2023.
- Louise, N. Why vertical AI agents could be 10x bigger than SaaS: Insights from Y Combinator, 2024. URL <https://techstartups.com/2024/12/03/why-vertical-ai-agents-could-be-10x-bigger-than-saas-insights-from-y-combinator/>.
- Luna-Ostaseski, G. Announcing the knowledge work demand index, 2021. URL <https://www.usebraintrust.com/blog/knowledge-work-demand-index-july-2021>.
- Martin, H. Salesforce will hire no more software engineers in 2025, says Marc Benioff, 2024. URL <https://www.salesforceben.com/salesforce-will-hire-no-more-software-engineers-in-2025-says-marc-benioff/>.
- McMahon Jr, M. J. and Abreu, A. G. Winner-take-all markets: Easing the case for progressive taxation. *Fla. Tax Rev.*, 4:1, 1998.
- Medaglia, R., Gil-Garcia, J. R., and Pardo, T. A. Artificial intelligence in government: Taking stock and moving forward. *Social Science Computer Review*, 41(1):123–140, 2023.
- Mollick, E. Prophecies of the flood, 2025. URL <https://substack.com/home/post/p-154445396>. One Useful Thing Substack newsletter post.
- Panait, C., Ljubenkov, D., and Alic, D. Striking the balance between innovation and regulation in ai-is europe leading the way or lagging behind. *Europuls Policy J. EU Affairs*, 1:27–45, 2021.
- Pifer, N. D., McLeod, C. M., Travis, W. J., and Castleberry, C. R. Who should sign a professional baseball contract? quantifying the financial opportunity costs of major league draftees. *Journal of Sports Economics*, 21(7):746–780, 2020.
- Qiao, D., Rui, H., and Xiong, Q. Ai and jobs: Has the inflection point arrived? evidence from an online labor platform. *arXiv preprint arXiv:2312.04180*, 2023.
- Rosen, S. The economics of superstars. *The American economic review*, 71(5):845–858, 1981.
- Saab, K., Tu, T., Weng, W.-H., Tanno, R., Stutz, D., Wulczyn, E., Zhang, F., Strother, T., Park, C., and Elahe Vedadi, e. a. Capabilities of gemini models in medicine, 2024.
- Salesforce. Introducing agentforce 2.0: The digital labor platform for building a limitless workforce, 2024. URL <https://www.salesforce.com/news/press-releases/2024/12/17/agentforce-2-0-announcement/>.
- Statista. Mean salary of MLB players from 2003 to 2023, 2024. URL <https://www.statista.com/statistics/236213/mean-salaray-of-players-in-majpr-league-baseball/>.
- U.S. Bureau of Labor Statistics. CPI inflation calculator, 2024. URL https://www.bls.gov/data/inflation_calculator.htm.
- U.S. Census Bureau. Acs population 25 years and over with a bachelor’s degree, 2023. URL <https://data.census.gov/table?q=b15012&y=2023>.
- Vale, R. Forecasting the 2022-23 tech layoffs using epidemiological models. *arXiv preprint arXiv:2305.05210*, 2023.
- Varanasi, L. Mark zuckerberg says ai could soon do the work of meta’s midlevel engineers, 2025. URL <https://www.msn.com/en-us/technology/artificial-intelligence/mark-zuckerberg-says-ai-could-soon-do-the-work-of-ar-BB1ri4RC>.
- Violante, G. L. Skill-biased technical change. *The new Palgrave dictionary of economics*, 2:1–6, 2008.
- Wang, K. What does law school cost, and is it worth it?, 2024. URL <https://www.forbes.com/advisor/education/law/law-school-cost/>.
- Web, H. B. Probability of playing college and professional baseball, 2024. URL <https://www.hsbbaseballweb.com/probability.htm>.
- Wiggers, K. Elon musk’s xai lands billions in new cash to fuel ai ambitions, 2024. URL <https://techcrunch.com/2024/12/25/elon-musks-xai-lands-billions-in-new-cash-to-fuel>.
- Witnauer, W. D., Rogers, R. G., and Saint Onge, J. M. Major league baseball career length in the 20th century. *Population research and policy review*, 26:371–386, 2007.

Woodhall, M. Earnings and education. In *Economics of Education*, pp. 209–217. Elsevier, 1987.

Zick, T., Kortz, M., Eaves, D., and Doshi-Velez, F. Ai procurement checklists: Revisiting implementation in the age of ai governance. *arXiv preprint arXiv:2404.14660*, 2024.