

KNOWLEDGE GUILDS: SHARING THE PRODUCTIVITY GAINS OF AI

PRELIMINARY & INCOMPLETE

*Salome Viljoen**, *Danielle Li***, *Zoe Cullen****

“What is happening macroeconomically is that distributed human labor is being used to train AI models by large, centralized actors, who concentrate the resulting profits while in the long run making the human labor obsolete ... An alternative that may work better for everyone, is to compensate data/content producers for their labor.”

- Dario Amodi, 2021¹

INTRODUCTION

AI models do not develop independently of the labor they may later be used to automate. Rather, the relationship between AI and labor is fundamentally extractive: workplace AI is typically trained on data generated by human work and, in many cases, derives its capacity to automate tasks precisely from its ability to disembody human knowledge. Customer service chat assistants are trained on transcripts of human support interactions.² Clinical decision support systems are trained on physicians’ notes, diagnostic codes, and treatment histories as captured by electronic health records.³ Logistics systems optimize routes using drivers’ historical movements and delivery decisions.⁴ Across many workplaces, labor itself generates *knowledge data*, records of workers’ expertise, which can then be used to train AI systems. When codified in this way, knowledge traditionally possessed by workers can become capital possessed by firms. We

* Assistant Professor, Michigan Law. The authors wish to thank [insert list of readers and editors here]

** Professor, MIT Sloan School of Management.

*** Associate Professor, Harvard Business School

¹ From unsealed documents in *Bartz v. Anthropic*. Dario Amodi is the CEO of Anthropic.

² See, e.g., CONSUMER FIN. PROT. BUREAU, *CHATBOTS IN CONSUMER FINANCE* 7 (2023), https://files.consumerfinance.gov/f/documents/cfpb_chatbot-issue-spotlight_2023-06.pdf.

³ See, e.g., Sriram Ragopal et al., *Artificial Intelligence-Based Clinical Decision Support in Pediatrics*, 93 *Pediatric Rsch.* 334, 336-37 (2023).

⁴ See, e.g., Matt Villano, *How Uber Freight is Leveraging AI to Make Truck Routes More Efficient*, MIT CENTER FOR TRANSPORTATION AND LOGISTICS (April 10, 2025), <https://ctl.mit.edu/news/how-uber-freight-leveraging-ai-make-truck-routes-more-efficient>; DHL FREIGHT, *How AI Improves Route Planning*, FREIGHT CONNECTIONS (July 15, 2025), <https://dhl-freight-connections.com/en/trends/ai-route-planning/>.

term this *knowledge capital*.

While the automation of work holds significant promise for productivity gains, it is not without risks for both workers and firms. By embedding human knowledge in scalable systems, AI models can extend work beyond the limits of an individual's time and physicality.⁵ The expertise of a top doctor, once codified in clinical software, can inform treatment decisions for patients across distance and over time. Whether that promise is realized, however, depends on how the risks and incentives surrounding AI development are managed. Workers may reasonably fear that their knowledge may be used to facilitate not only their own obsolescence, but that of others in their occupation. Firms, in turn, may worry that such fears will lead workers to withhold knowledge, resulting in incomplete data and lower-quality AI. If knowledge data is mismanaged, the result may be inferior AI systems, diminished productivity, and more exploitative labor relations. But a better legal framework for governing knowledge data can align incentives and fairly distribute both risks and rewards, they can facilitate innovation that benefits both workers and firms.

This Article makes three contributions. In Part I, it defines the concept of knowledge data and knowledge capital. This part explains how expanded workplace surveillance and AI tool use have made it possible to codify increasingly broad forms of worker knowledge as AI, including knowledge that was formerly tacit and embedded within individual workers. It also links the growing worker backlash against AI to longer historical trends of automation as a means of disempowering workers. In Part II, the Article develops a framework for understanding how the creation of AI capital from worker knowledge affects different workers in different ways. Part II introduces a taxonomy based on two key dimensions: replicability—how easily a worker's knowledge can be modeled by an employer—and substitutability—how easily that knowledge can be

⁵Whether and to what extent AI can substitute for human expertise remains the subject of ongoing empirical research and debate. We do not take a position in that debate and indeed evidence is mixed. For example, see e.g. the Carnegie Mellon agent software company: <https://tech.yahoo.com/ai/articles/next-assignment-babysitting-ai-081502817.html> Instead, for purposes of analysis, we assume that meaningful substitution is or will become possible in at least some contexts, in order to examine its potential workplace effects. Regardless of whether this assumption is ultimately borne out, legislators and legal scholars must plan for the social and economic disruptions that more capable AI systems may produce. AI need not fully or perfectly replace human labor to generate significant disruption.

supplied by others. Under current institutional arrangements, many workers across this taxonomy (and the firms that employ them) are likely to experience inequitable or inefficient outcomes. In Part III, the Article proposes data intermediaries, which we call *knowledge guilds*, as a potential legal response. Knowledge guilds are a legal template for collective knowledge data governance. They can be achieved via a variety of forms, but knowledge guilds must combine two essential elements. First, they entitle workers to data sovereignty over their knowledge. Data sovereignty recognizes that workers have three equal interests at stake in knowledge data: *control* over how and for what data is used, the financial *value* data helps produce, and the *privacy* risks data use may create. Second, knowledge guilds are a collective bargaining mechanism. Workers need to bargain over and enforce claims to knowledge data collectively if they are to exert meaningful power onto firms.

To understand the relationship between labor, knowledge data, and knowledge capital more concretely, consider the following two examples.

- *Customer Service*: AI chat assistants in call centers are trained on transcripts of thousands of recorded customer interactions, capturing how agents resolve problems, de-escalate frustrated customers, and diagnose technical issues. These recordings function as knowledge data and are typically collected passively through surveillance systems (“calls are recorded for quality assurance”). This knowledge base is not static. As firms introduce new products, update policies, or encounter novel customer problems, agents generate fresh transcripts that reflect new solutions. To build high-performing models, firms often identify top agents and tag their conversations as exemplars during model development. The resulting AI chat assistant can thus be understood as a form of knowledge capital: it codifies the expertise and interpersonal style of specific workers into an ownable, scalable, and transferable asset that exists independently of the individuals who produced it.
- *Consulting*: AI tools that support consultants are typically built from materials already generated and stored within firms’ digital systems: slide decks and strategy memos created in enterprise Microsoft 365 environments, emails sent through monitored corporate accounts, and project histories logged in collaboration platforms such as Slack, Asana, or Microsoft Teams. Yet unlike chat transcripts, knowledge data relevant for consulting are less structured and, at least currently, the process of customizing an AI assistant for knowledge work generally

requires workers themselves to curate the information they provide to AI models. Because of this, consultants have greater discretion over the amount and quality of knowledge data they provide. Effective models require continual updates reflecting new client engagements and shifting market or economic conditions, meaning their performance depends on fresh inputs that consultants themselves choose how fully to document and share.

As these examples demonstrate, workers in modern organizations create two streams of value for the firm.⁶ First, they work: they help a customer or write a report. This is labor as classically conceived: the contribution of human time and skill to production, for which workers are paid a wage. Second, they produce knowledge data. Whether transcripts, recordings, or documents, these records serve as examples that can potentially be used to train AI models to replicate that same work.⁷ To understand the distinct value and legal import of this data, it is helpful to distinguish its production and its value from the underlying labor tasks that this data captures in digital form. The extraction and use of knowledge data is broadly sanctioned within the scope of existing employment contracts, even as its use to train AI modes generated increased worker and social backlash.

Both examples illustrate the productivity potential of AI systems. A recent study by Brynjolfsson, Li, and Raymond shows that the introduction of AI chat assistants improved productivity by giving all workers in the firm access to the problem solving skills codified in the knowledge data of top performers.⁸ Similarly, Dell'Acqua et. al. shows that management consultants with access to AI tools improved the speed and quality of writing and ideation tasks.⁹

⁶ This applies to modern firms broadly, and not just firms that have introduced AI, because the only necessary condition is that the firm needs to have surveillance so that their labor data *could* be used to build AI systems -- these AI systems could be used at the focal firm or not. On data value, see generally Amanda Parsons & Salomé Viljoen, *Valuing Social Data*, 124 COLUM. L. REV. 993 (2024).

⁷ SHOSHANA ZUBOFF, *IN THE AGE OF THE SMART MACHINE: THE FUTURE OF WORK AND POWER* (1988). Zuboff describes the investments and changes to workflows needed to make tasks more legible via data flows as 'informatting'. 'Datafication' is another common way to refer to the set of practices and techniques involved in making labor legible as data.

⁸ Erik Brynjolfsson, Danielle Li & Lindsey Raymond, *Generative AI at Work*, 140 Q.J. ECON. 889, 890, 900 (2025).

⁹ Fabrizio Dell'Acqua et al., *Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality* 17 (Harv. Bus. Sch. Working

At the same time, however, the transformation of knowledge into knowledge capital raises pressing concerns over fairness and exploitation. Whether call center transcripts, code repositories, or digital art, AI models are often trained on workers' data without compensation nor fully informed consent.¹⁰ In the Brynjolfsson, Li, and Raymond call center study, access to AI tools increased the problem resolution rates of novice workers by over 30%, but did not improve performance among top performers.¹¹ In effect, the knowledge data of top performers was used to improve firm profits, as well as the wages and productivity of others, without benefit to them. If anything, the availability of the AI model allows the firm to hire less-skilled workers into the role, potentially at lower wages.¹² In the longer run, many workers fear that the AI systems they are helping to train may replace them entirely.¹³

In light of these dynamics, workers may have reason to resist efforts to codify their expertise through AI systems.¹⁴ Although AI adoption may free some workers to pursue more interesting or better-paid tasks, it may also increase the risk of replacement. Faced with that uncertainty, workers may have few incentives to accelerate a process that could diminish their own job security. In highly surveilled and data-legible roles such as remote customer support,

Paper, No. 24-013, 2023), <https://www.hbs.edu/faculty/Pages/item.aspx?num=64700>.

¹⁰ See Frank Pasquale & Haochen Sun, *Consent and Compensation: Resolving Generative AI's Copyright Crisis*, 110 VA. L. REV. ONLINE 207, 209-10 (2024).

¹¹ Brynjolfsson et al., *supra* note X, at 891.

¹² See generally HAGEN BLIX & INGEBORG GLIMMER, *WHY WE FEAR AI: ON THE INTERPRETATION OF NIGHTMARES* (2025) (arguing AI is not a productivity enhancing tool but instead a wage suppression tool).

¹³ See Lakshmi Varanasi, *Vercel trained an AI agent on its best salesperson. Then it cut the 10-person team down to 1*, BUS. INSIDER, (Oct. 27, 2025), <https://www.businessinsider.com/ai-agent-entry-level-sales-jobs-vercel-2025-10>; EDELMAN, 2026 EDELMAN TRUST BAROMETER 12 (2026), https://www.edelman.com/sites/g/files/aatuss191/files/2026-01/2026%20Edelman%20Trust%20Barometer%20Global%20Report_Final.pdf (Showing two-thirds of low-income U.S. survey respondents said that “people like me will be left behind rather than realize any real advantages from generative A.I.” Perhaps more remarkably, nearly half of high-income respondents felt the same way.).

¹⁴ Indeed, concerns over AI use and job security have been at the heart of several recent labor disputes. See Molly Kinder, *Hollywood writers went on strike to protect their livelihoods from generative AI. Their remarkable victory matters for all workers*, BROOKINGS (April 12, 2024), <https://www.brookings.edu/articles/hollywood-writers-went-on-strike-to-protect-their-livelihoods-from-generative-ai-their-remarkable-victory-matters-for-all-workers/>; Tom Richardson, *Video Game Actors' Strike Officially Ends After AI Deal*, BBC NEWSBEAT (July 10, 2025), <https://www.bbc.com/news/articles/c5ykx117keqo>.

workers may have limited ability to prevent the collection of their knowledge data. For the vast majority of occupations, however, building AI models that fully capture human expertise is difficult, especially without a worker's cooperation. Even when much of a process and its outputs are recorded, gaps remain between what a firm can observe and the full judgment workers bring to their role. Consultants, doctors, salespeople, and software developers all retain substantial discretion over how thoroughly they document their processes or articulate their reasoning. When AI systems are viewed as adversarial, individuals may degrade the quality of knowledge data they supply by simply failing to do their best to make their tacit knowledge explicit.

Such actions can be costly. When workers are not incentivized to share high-quality knowledge data, AI systems are more likely to be built on incomplete or degraded inputs, producing weaker models that fail to capture genuine expertise. If these models are nonetheless cost-effective, firms may deploy them in place of more expensive human experts. In that scenario, firms reduce costs, but society receives lower-quality service and skilled workers face exploitation and displacement.

Knowledge guilds offer a promising legal solution. In particular, knowledge guilds can accomplish three valuable objectives regarding AI implementation and development. First, by enabling the conditions of worker cooperation, they can facilitate and encourage high quality AI innovation. Second, they can ensure that workers whose expertise are inputs into a model get their share of the value they helped create, and who may suffer career losses as a result of labor automation. Third, knowledge guilds can serve to legitimize and empower mechanisms for people more broadly to recognize the value their data provides, and provide one means by which to empower data providers against its exploitation. Although they cannot resolve every uncertainty or legal harm associated with workplace AI, properly designed knowledge guilds can nonetheless meaningfully reduce incentive misalignment and mitigate important sources of harm.

I. KNOWLEDGE DATA AND ITS ROLE IN AI DEVELOPMENT

Part I introduces the notion of labor-data: what it is, grounding the concept in prior literature on workplace automation, and introducing the law's role in managing and mediating how value from workplace innovation is shared across

different settings and between workers and firms. Part I will review prior work on labor as data and situate that work in contemporary debates about the role of AI in shaping the present and future of work. Part I will lay out the social challenges and labor risks that arise from the status quo approach to datafying labor for AI automation].

A. *What is Knowledge Data*

Workers in modern workplaces are increasingly subject to surveillance technologies—audio, video, and screen recording, as well as the monitoring of workplace communications. Through the act of working, increasingly in tandem with AI systems responsive to feedback, these systems generate a continuous stream of data about work: records of processes, decisions, and outputs. We refer to this as labor data. Labor data mediates the transformation of labor into capital by extracting traces of embodied know-how and embedding them in software systems that guide or automate future work.

The capacity to scale human expertise beyond the limits of individual workers holds enormous potential for productivity and consumer welfare. While the record is mixed, some evidence of such gains is already emerging. In customer support, Brynjolfsson, Li, and Raymond study the staggered roll-out of a generative-AI assistant to 5,172 agents and find a 15 percent average increase in issues resolved per hour, with the largest gains for novices (\approx 34 percent).¹⁵ Trained on the transcripts and behavioral data of top-performing agents, the system disseminated their (formerly) tacit skills across the organization, effectively spreading expertise and raising overall output.¹⁶

Similar dynamics appear in higher-skill settings. Cui et al. conduct a large-scale field experiment randomly assigning access to GitHub Copilot for 4,867 developers at Microsoft, Accenture, and a Fortune 100 manufacturer.¹⁷ Copilot access led to a 25 percent increase in completed tasks, with the largest gains among recent hires and junior developers, again consistent with AI diffusing expert patterns to less-experienced workers. Beyond these studies, one can readily imagine other domains where scaling human expertise could yield comparable benefits—for example, extending the diagnostic reach of top physicians to patients in rural or resource-limited settings.

¹⁵ Brynjolfsson et al., *supra* note X, at 890-91.

¹⁶ *Id.*

¹⁷ Kevin Zheyuan Cui et al., *The Effects of Generative AI on High-Skilled Work: Evidence from Three Field Experiments with Software Developers*, MGMT. SCI. (forthcoming 2026), <https://doi.org/10.1287/mnsc.2025.00535>.

At the same time, use of labor data to build AI capital raises concerns about the potentially uncompensated and uncontrolled use of worker-generated data. In media and entertainment, for example, the SAG-AFTRA and Writers Guild of America strikes of 2023–25 spotlighted the risks when workers’ voices, likenesses, or performance logs are harvested without meaningful consent, compensation, or control.¹⁸ The 2024–25 SAG-AFTRA video-game strike, for instance, centered on studios’ creation of digital replicas of performers’ voices and motion-capture data that could be reused indefinitely without additional pay.¹⁹ Beyond entertainment, workers across industries face opaque logging of their interactions, keystrokes, and screen data, which then feed AI systems in ways they may neither consent to nor benefit from.²⁰

Situations where workers do not benefit from AI models built from their labor data—and indeed, where they may be negatively impacted by such models—generate concerns about both fairness and efficiency. If workers feel exploited, they may resist participation in data-capture regimes that are crucial to the development of future AI models.

Cullen, Li, and Li formalize this problem, showing that workers’ willingness to supply labor data depends critically on how they believe that data will be used.²¹ In a large online experiment, they randomly informed workers about how data from common forms of workplace surveillance—such as audio or screen recording—can be used to train AI systems that perform similar tasks, like customer communication or expense reporting.²² Workers who received this information became less willing to share data with their employers, both in self-reports about their primary jobs (expressing reduced willingness to document knowledge and workflows) and in incentive-compatible tests related to their

¹⁸ See Kinder, *supra* note X.

¹⁹ See Richardson, *supra* note X.

²⁰ James D. Walsh, ‘*There’s Just No Reason to Deal With Young Employees’ AI is taking entry-level jobs. What happens when Gen-Z-ers can’t start their careers?*’, N.Y. Mag. (Nov. 12, 2025), <https://nymag.com/intelligencer/article/ai-replacing-entry-level-jobs-gen-z-careers.html>, (Interviewing a recent college graduate employed as a copywriter. “For the first six months of the job, AI was hardly ever mentioned. Then her bosses began encouraging her to “experiment” with ChatGPT. Last spring, AI became a requirement of the job. ... ‘It’s so obvious I’m training something to replace me, and I think that’s kind of why I’ve been hesitant. Because if I can make it do my job successfully, they don’t need me to do my job anymore.”).

²¹ Zoë Cullen, Danielle Li & Shengwu Li, Labor as Capital: AI and the Ownership of Expertise 1-3 (Oct. 20, 2025) (preliminary & incomplete draft), https://zcullen.github.io/assets/docs/Labor-as-Capital_WP.pdf.

²² *Id.*

secondary work as survey takers (declining substantial payments for additional data about their preferences).

When workers expect that their contributions will strengthen employer-owned models that could replace them, they have incentives to withhold or distort behavior, reducing the informational value of their data. This dynamic is inefficient for both sides: firms lose the ability to train accurate models, and workers forgo both potential productivity gains in their work and potential shares in the wealth generated from greater productivity. The analysis further shows that data ownership shapes these incentives. When workers own their labor data, they regain motivation to contribute it, but ownership structure determines who benefits. Under individual ownership, workers compete to sell similar data, driving prices down even when labor data is highly valuable as a whole. Collective data control, by contrast, can sustain incentives for high-quality data generation while ensuring that workers as a group capture a meaningful share of the gains. Ownership and governance of labor data are therefore not only fairness concerns but central determinants of economic efficiency.

Automation that unsettles the relationship between labor and capital is not new. Indeed, automation becoming a flashpoint where labor and capital to renegotiate their respective shares from production is, in some sense, the story of industrial capitalist innovation and struggle.²³ AI is in many ways just the latest technology that aims to automate a new class of cognitive human expertise.

Nor is the importance of observing workers to the promise of automation as the path to productivity enhancement new. Would-be automators have long observed workers to inform automation design. Taylor pioneered time studies and the development of scientific management premised on close observations of workers at their tasks.²⁴ Frank and Lillian Gilbreth built on these in their time and motion studies, and Ford implemented and built on Taylorist practices to

²³ On role of law in mediating these struggles, see work of Amy Kapczynski, Yochai Benkler. See also Julie Cohen on information law; Katharina Pistor.

²⁴ See generally FREDERICK WINSLOW TAYLOR, *THE PRINCIPLES OF SCIENTIFIC MANAGEMENT* (1911); *Testimony of Frederick Winslow Taylor at Special Hearings Before Special Committee of the House of Representatives, January 1912*, 11 BULL. TAYLOR SOC'Y 95 (1926), <https://hdl.handle.net/2027/mdp.39015046810142>; HUGH G. J. AITKEN, *SCIENTIFIC MANAGEMENT IN ACTION: TAYLORISM AT WATERTOWN ARSENAL, 1908–1915*(1960); ROBERT KANIGEL, *THE ONE BEST WAY: FREDERICK WINSLOW TAYLOR AND THE ENIGMA OF EFFICIENCY* (1997).

great effect at his Dearborn factories.²⁵ The promise of labor data for automation was not lost on these early management scientists. Indeed, the Gilbreth's noted of their studies that "in principle, every relevant detail should be captured and subject to investigation and optimization".²⁶

This history also shows the fractious relationship between labor and capital, who then, as now, sharply diverge on automation.²⁷ During his time, organized labor viewed Taylor's optimization methods with skepticism (to put it lightly). Implementation of scientific management practices precipitated work stoppages and other forms of worker organizing.²⁸ On their account, management science was as interested, if not more, in weakening labor's control over the workplace (and labor's share in the fruits of production) as it was in enhancing overall productivity.²⁹ And indeed, Taylor believed in transferring control from "autonomous and inefficient" workers to their managers. In their critique of Taylorism, laborers urged a Veblenian account of automation. On this account, capitalist investment in innovation is agnostic about whether it produces increased returns from productivity gains or from increasing capital's share of existing production. The incentive to automate is not a disinterested pursuit for productivity gains, nor is its overall social effect separable from the background distributions of rights and obligations that capital and labor owe one another.³⁰

²⁵ FRANK GILBRETH, MOTION STUDY (1911); FRANK GILBRETH & LILLIAN GILBRETH, APPLIED MOTION STUDY (1917); FRANK GILBRETH & LILLIAN GILBRETH, FATIGUE STUDY (1916); STEPHEN MEYER III, THE FIVE DOLLAR DAY: LABOR MANAGEMENT AND SOCIAL CONTROL IN THE FORD MOTOR COMPANY 1908-1921, AT 11-36 (1981).

²⁶ Marion Fourcade & Kieran Healy, *The Ordinal Society* 73 (2024); *see also* James Rinehart, *The Tyranny of Work* (1975).

²⁷ Colleen McClain et al., Pew Rsch. Ctr., *How the U.S. Public and AI Experts View Artificial Intelligence* 6 (2025); Luona Lin & Kim Parker, Pew Rsch. Ctr., *U.S. Workers Are More Worried Than Hopeful About Future AI Use in the Workplace* 14-17 (2025).

²⁸ Mike Davis, *The Stopwatch and the Wooden Shoe: Scientific Management and the Industrial Workers of the World*, *Radical Am.*, Jan.-Feb. 1975, at 69, 73-74, <https://library.brown.edu/pdfs/1125403651398134.pdf>.

²⁹ *Id.* at 69-73..

³⁰ Labor unions and others also contested the claim that automation always resulted in / was pursued for productivity gains. Sometimes, they asserted, innovation was pursued that merely reallocated the surplus of production from labor to capital. *See, e.g.*, Marc-André Gagnon, *Capital, Power and Knowledge According to Thorstein Veblen: Reinterpreting the Knowledge-Based Economy*, 41 *J. Econ. Issues* 593, 597 (2007), <http://www.jstor.org/stable/25511213>. As Landon Winner notes, "If we suppose that new technologies are introduced to achieve increased efficiency, the history of technology shows that we will sometimes be disappointed. Technological change expresses a panoply of human motives, not the least of which is the desire of some to have dominion over others even though it may require an occasional sacrifice of cost savings and some violation of the normal standard of trying to get more from less." *The Whale and the*

These past accounts echo in contemporary debates over AI's effect on the workplace. Some argue AI is operating as a tool of wage suppression.³¹ And there is some emerging evidence of worker job loss, which both CEOs and other observers have linked with AI.³² Unsurprisingly, evidence of worker frustration at AI is growing. In film, the introduction of Tilly Norwood — a fully AI-generated "actress" trained on performances of real actors without consent or compensation — provoked immediate backlash from SAG-AFTRA, individual actors, and multiple talent agencies, who saw it as a direct threat to performers' livelihoods and a violation of their data sovereignty.³³ Prominent filmmakers have taken strong personal stands: Guillermo del Toro declared he would "rather die" than use generative AI, comparing tech industry leaders to Mary Shelley's Victor Frankenstein.³⁴ In music, Paul McCartney contributed a near-silent track to *Is This What We Want?*, a compilation by over 1,000 musicians protesting the UK government's proposed copyright changes that would allow AI developers to train on creative works unless artists affirmatively opted out.³⁵ And at the other end of the labor market, Kenyan data labelers — the workers who annotate, label, and curate the training data that makes AI systems function — formed the Data Labelers Association in early 2025 to challenge wages that amount to pennies per task and working conditions they describe as systematically unjust.³⁶ Unsurprisingly then, we also see increased evidence of

Reactor 24 (1986).

³¹ Brian Merchant, "AI is an attack from above on wages": An interview with cognitive scientist Hagen Blix, *Blood in the Mach*. (October 1, 2025), <https://www.bloodinthemachine.com/p/ai-is-an-attack-from-above-on-wages>.

³² AI can suppress hiring or wage growth both due to the opportunity costs of investing firm resources in AI instead of in workers, and in the direct replacement due to the ability of AI to replace labor. Evidence suggests that some of both is happening. Take the recent example of Amazon announcing tens of thousands of job cuts. Amazon SVP Beth Galetti wrote in a memo accompanying the layoff announcement that "AI is the most transformative technology we've seen since the Internet" and as such Amazon needs "to be organized more leanly, with fewer layers and more ownership." In other words, AI is driving the cost-cutting, perhaps hitting both of the above buckets at once.

³³ Chloe Veltman, *Could 'the Next Scarlett Johansson or Natalie Portman' Be an AI Avatar?*, NPR (Oct. 1, 2025), <https://www.npr.org/2025/09/30/nx-s1-5558119/tilly-norwood-ai-hollywood>.

³⁴ Terry Gross, *Filmmaker Guillermo del Toro Says 'I'd Rather Die' Than Use Generative AI*, NPR (Oct. 23, 2025), <https://www.npr.org/2025/10/23/nx-s1-5577963/guillermo-del-toro-frankenstein>.

³⁵ David Browne, *Paul McCartney Protests AI with Silent Track*, *Rolling Stone* (Nov. 17, 2025), <https://www.rollingstone.com/music/music-news/paul-mccartney-ai-silent-track-is-this-what-we-want-1235466449/>.

³⁶ Sebastian Klovig Skelton, *Kenyan AI workers form Data Labelers Association*, *Comput. Wkly*. (Feb. 14, 2025), <https://www.computerweekly.com/news/366619321/Kenyan-AI-workers->

worker frustration at AI, though we have yet to see many mass walkouts. Beyond work stoppages and protests, 2025 saw the rise of novel "techno-refusal" tactics: artists using "data poisoning" tools (like Nightshade and Glaze) to corrupt training data; and web developers deploying "tarptitting" — trapping AI crawlers in endless loops to prevent content scraping.³⁷

Beyond worker frustration, research on the actual roll-out of AI in firms reveals that organizations struggle to translate AI training into meaningful adoption by employees. Dillon, Jaffe, Immorlica, and Stanton report results from a six-month, cross-industry randomized field experiment providing 7,137 knowledge workers access to a generative AI tool integrated into their existing email, document, and meeting applications.³⁸ They find that workers did not significantly change time spent in meetings or restructure their broader task composition. Usage also declined from its initial peak as the novelty wore off.³⁹ Humlum and Vestergaard link large-scale adoption surveys to administrative labor records in Denmark, covering 25,000 workers and 7,000 workplaces across eleven AI-exposed occupations.⁴⁰ Despite widespread firm investments in training initiatives the authors estimate precise null effects on recorded work hours and earnings. Together, these studies suggest that even when firms attempt to promote AI adoption, employee take-up may fall short of its potential — echoing Acemoglu and Johnson's broader warning that automation can yield disappointing productivity outcomes if misaligned with employee objectives.⁴¹

The practice of surveilling workers to inform capital innovation is not new. Frederick W. Taylor pioneered time studies and the development of scientific management over a century ago from close observations of workers at their

form-Data-Labelers-Association.

³⁷ Melissa Heikkilä, *The AI lab waging a guerrilla war over exploitative AI*, MIT Tech. Rev. (Nov. 13, 2024), <https://www.technologyreview.com/2024/11/13/1106837/ai-data-positoning-nightshade-glaze-art-university-of-chicago-exploitation/> (data poisoning tools); Ashley Belanger, *AI haters build tarpits to trap and trick AI scrapers that ignore robots.txt*, Ars Technica (Jan. 28, 2025), <https://arstechnica.com/tech-policy/2025/01/ai-haters-build-tarpits-to-trap-and-trick-ai-scrapers-that-ignore-robots-txt/> (tarptitting and data poisoning).

³⁸ Eleanor W. Dillon et al., *Shifting Work Patterns with Generative AI 1* (NBER Working Paper 33795, Nov. 2025), <https://doi.org/10.3386/w33795>.

³⁹ *Id.*

⁴⁰ Anders Humlum & Emilie Vestergaard, *Large Language Models, Small Labor Market Effects 1* (NBER Working Paper 33777, Oct. 2025), <https://doi.org/10.3386/w33777>.

⁴¹ See generally DARON ACEMOGLU & SIMON JOHNSON, POWER AND PROGRESS (2023). [ntd: cite to additional industry surveys on this point.]

tasks. Taylor was explicit about the direction of knowledge transfer: managers, he wrote, should assume "the burden of gathering together all of the traditional knowledge which in the past has been possessed by the workmen and then of classifying, tabulating, and reducing this knowledge to rules, laws, and formulæ."⁴² Frank and Lillian Gilbreth extended this logic to its granular extreme, deploying motion-picture cameras to decompose every worker's task into elemental units of motion—eighteen distinct categories they called "therbligs"—filmed at fractions of a second, creating a permanent visual record from which the "one best way" to perform any job could be extracted, codified, and replicated without the originating worker's further input. Harry Braverman identified the core dynamic at work: Taylorism was premised on the expropriation of workers' craft knowledge, its concentration in the hands of management, and the use of this monopoly of knowledge as power over workers—what he termed the "separation of conception from execution."⁴³ Taylor himself saw no need for unions once this transfer was complete, claiming that scientific management would render "labor unions and strikes unnecessary."⁴⁴

This history also reveals a recurrent pattern of worker resistance. The introduction of time studies at the government-owned Watertown Arsenal in 1911 provoked a walkout by foundry workers and a Congressional investigation in which legislators concluded that Taylor's system was "arbitrary and harsh" and "injurious to the worker's manhood and welfare," ultimately banning stopwatch studies and bonus systems in government establishments. The American Federation of Labor organized national opposition, and the conflict contributed to the crystallization of the modern industrial union as a countervailing institution.

A later episode sharpens the parallel with contemporary AI. In the 1980s, the dominant paradigm in artificial intelligence was the "expert system"—a rule-based program designed to replicate human expert judgment by encoding an expert's reasoning into if-then rules. Two-thirds of Fortune 500 companies adopted the technology, and Edward Feigenbaum's *The Rise of the Expert Company* promised that captured expertise could be replicated and scaled across

⁴² Talyor, *supra* note X, at 36.

⁴³ HARRY BRAVERMAN, *LABOR AND MONOPOLY CAPITAL* 79 (25th anniversary ed. 1998).

⁴⁴ CTK

entire organizations.⁴⁵ The process required "knowledge engineers" to interview human experts at length, extracting and codifying their tacit reasoning. But the approach ran aground on what researchers called the "knowledge acquisition bottleneck": experts often could not articulate what they knew, and—crucially—many actively sabotaged the process, recognizing that encoding their expertise into a machine meant rendering themselves dispensable. Expert systems projects were, in the words of one practitioner, "often sabotaged by the experts themselves, who pretended not to be available... 'Am I supposed to give this damn system MY knowledge so I can be replaced by a computer? No way!'"⁴⁶ The one cooperative expert in a widely cited case study was notable precisely because he was approaching retirement and had no stake in protecting his position.⁴⁷

Where expert systems required conscious, cooperative participation from the expert, large language models and generative AI systems train on the accumulated digital residue of millions of workers' outputs: their texts, code, images, and performances, harvested at scale and often without explicit consent. The tacit knowledge that once served as a worker's most durable bargaining chip—the expertise that could not be extracted without the worker's active cooperation—now leaks through the digital traces of ordinary work. In this light, the contemporary AI moment represents not a break from the Taylorist project but its fulfillment: the complete datafication of worker knowledge, achieved not through the stopwatch or the knowledge engineer's interview, but through the ambient capture of labor as training data.

B. *The Worker-Firm Knowledge Gap and Suboptimal AI Outcomes*

1. Defining the worker–firm knowledge gap

Despite widespread experimentation with generative AI tools, many firms report that these technologies have yet to deliver transformative productivity gains.⁴⁸ One important reason is that, even when AI systems have access to vast

⁴⁵ See generally EDWARD FEIGENBAUM, PAMELA MCCORDUCK & H. PENNY NII, *THE RISE OF THE EXPERT COMPANY* (1988).

⁴⁶ Rafe Brena, *What Happened with Expert Systems?*, TOWARDS DATA SCI. (Apr. 27, 2024), <https://towardsdatascience.com/what-happened-with-expert-systems-aad399cab180/>.

⁴⁷ *Id.*

⁴⁸ See discussion in Part IA *supra*. There are several ongoing surveys of firms reporting disappointment or non-existent productivity gains. <https://fortune.com/2026/02/17/ai-productivity-paradox-ceo-study-robert-solow-information-technology-age/>; https://mlq.ai/media/quarterly_decks/v0.1_State_of_AI_in_Business_2025_Report.pdf

amounts of formal knowledge and computational power, they often lack the contextual and tacit expertise held by workers who perform tasks on a daily basis. Workers frequently possess local monopolies over the various kinds of information relevant for performing a task successfully.

Contextual knowledge includes details such as the preferences of a particular client, the informal norms of a workplace, or the most effective ways to deviate from standard operating procedures to get work done in practice. This knowledge is highly specific to time, place, and organizational setting, and is therefore rarely documented or incorporated into the data used to train foundational AI models. In addition to context, workers often possess tacit knowledge about how to perform tasks. Tacit knowledge describes skills and judgments that are not codified or formalized, often because those who possess them cannot fully articulate how they operate. By definition, such knowledge has historically resisted easy datafication. In practice, workers' expertise typically combines both forms: an intuitive sense of what will work in a particular setting and what will not.

For example, a call center worker's effectiveness in handling an angry customer depends not only on scripted best practices learned in training, but on an accumulated, largely unconscious ability to read cues and adjust tone, pacing, and language in real time—often informed by prior interactions with similar customers or with this specific caller. A trauma surgeon may likewise “know,” in a way that resists precise articulation, how much pressure to apply to a severed artery, drawing on embodied experience that integrates sight, touch, and situational urgency. A product manager may sense that a particular feature, though defensible on paper, will provoke resistance from a specific client stakeholder or internal champion, based on subtle organizational dynamics and past interactions that are never formally recorded.

Together, these examples highlight that the worker-firm knowledge gap encompasses emotional, embodied, and intuitive forms of expertise that are difficult to observe, codify, or transfer.

2. Why might such a gap exist?

In a growing number of workplace settings, modern surveillance tools are capable of capturing aspects of workers' tacit or contextual knowledge, thereby narrowing the worker-firm knowledge gap. For example, customer service calls

can be recorded, transcribed, and analyzed at scale, allowing firms to train AI systems that reproduce some of the emotional cues, phrasing, and timing used by highly effective agents.

In many other cases, however, codifying what workers know remains very difficult. At a technical level, much workplace knowledge is *illegible* in the sense that it does not readily translate into machine-readable data. For instance, while customer interactions can be recorded, it is far more difficult to capture a scientist's tacit reasoning or contextual judgment, because we lack reliable tools for observing and recording the process of thinking itself. Even when data capture is technically feasible, organizational and incentive constraints may prevent firms from accessing workers' knowledge. Workers may actively resist providing data that would erode their expertise or threaten their job security, as in the labor conflicts discussed above. Alternatively, workers may simply be under-incentivized to generate detailed, legible records of their work, which takes effort and time away from other tasks. A physician, for example, may prioritize patient care over producing exhaustive clinical notes, even if those notes would be valuable inputs for training AI systems. As a result, the worker-firm knowledge gap can persist even in environments with advanced data collection technologies.

Consequences of the worker-firm knowledge gap

The worker-firm knowledge gap provides a useful lens for understanding why AI systems may fall short of the performance levels desired by workers, firms, or society. We focus on three related forms of suboptimal outcomes that can arise from this gap: the deployment of AI systems that function, but at degraded quality; the exploitation of workers as a means of extracting the knowledge needed to improve AI systems; and the failure to adopt potentially valuable AI systems when workers take steps to avoid such exploitation.

Before proceeding, we want to make clear that this analysis is not intended to provide a comprehensive accounting of the welfare effects of AI adoption. Even when a worker-firm knowledge gap leads to lower-quality AI outputs or products, adoption may still be socially beneficial if it substantially reduces costs, either by increasing firm profitability or by lowering prices for consumers. Such cost reductions can allow firms to remain viable, supporting greater market participation, product variety, and competition. At the same time, lower prices may encourage higher levels of consumption, which can generate negative

externalities, such as environmental degradation, that are not reflected in market prices. The overall welfare impacts of suboptimal AI will depend not only on the magnitude of any resulting price decreases, but also on the scale of any accompanying externalities.

a. A) Adoption of degraded AI systems

AI systems are more likely to underperform when they lack access to workers' tacit and contextual expertise. In customer service, this can produce interactions that feel rigid or inauthentic, as systems apply uniform displays of "corporate empathy" across customers rather than inferring (as skilled agents do) when a more direct approach would be more effective. In clinical settings, decision-support tools may rely heavily on population-level evidence from clinical studies while overlooking information that physicians routinely incorporate, such as a patient's lifestyle, preferences, or cultural context. As a result, recommendations may be technically sound yet poorly matched to individual patients, reducing trust, adherence, and overall effectiveness. Similarly, in organizational decision-making, AI tools trained primarily on formal guidelines and policies may ignore informal norms or interpersonal dynamics, leading to recommendations that workers view as unrealistic or impractical.

Despite these quality concerns, firms may still rationally adopt AI technologies that perform worse than human labor. When AI enables substantial reductions in labor costs, the resulting savings can outweigh revenue losses from degraded product or service quality. Consider self-checkout systems. Consumers frequently report frustration with self-checkout technologies that are overly sensitive or unreliable, and firms report higher rates of customer theft, both indicators of diminished service quality. Yet the labor cost savings are often substantial enough that firms have little incentive to abandon or significantly redesign these technologies.

Over longer horizons, repeated cost-cutting through automation can reshape production itself in ways that further erode quality. The history of garment manufacturing provides a clear example: successive waves of automation and outsourcing progressively de-skilled labor—from bespoke tailoring to factory-based fast fashion—alongside a sustained decline in clothing quality. In this sense, short-run adoption decisions driven by labor cost savings

can accumulate into long-run changes in product quality and production structure.

Market structure further shapes these incentives. In markets with limited competition, firms may face few consequences for degrading product or service quality. In more competitive markets, firms may still have incentives to maintain quality. Yet, perversely, competition can also have the opposite effect: intense price pressure may push firms to adopt cost-cutting AI technologies even more aggressively, particularly when lower costs allow firms to compete more aggressively on prices.

b. B) Exploitation of workers

Our definition of “bad AI” encompasses not only the performance of AI systems, but also the ethical conditions under which those systems are developed. When firms recognize that workers possess tacit and contextual expertise that is valuable for improving AI systems, their response need not be to create collaborative arrangements that encourage workers to share that knowledge. Instead, firms may attempt to narrow the worker–firm knowledge gap by investing in increasingly invasive forms of monitoring and surveillance that extract, codify, and appropriate workers’ expertise without meaningful consent or compensation.

For example, production studios have invested in body scanners that capture performers’ biometric data and likeness. Such technologies make it possible for studios to create a digital replica of a performer that could, in theory, be used in lieu of re-shooting scenes, etc. Actors have reported being asked to submit to body-scanning, often without clear information about how that data will be used or what rights they retain over it.⁴⁹ Many performers feel pressured to comply for fear of seeming difficult and jeopardizing their careers, even as the body scan data they are being asked to provide could itself jeopardize their careers.

These dynamics arise in part because advances in digital sensing and surveillance make it increasingly feasible for firms to observe and record

⁴⁹ Michael Savage, *Have We Done Ourselves Out of a Job?: Concerns in Film and TV Industry over On-Set Body Scanning*, Guardian (Oct. 17, 2025), <https://www.theguardian.com/film/2025/oct/17/ai-data-scanning-film-tv-actors-crew>.

dimensions of work that were previously informal, private, or difficult to measure. Tools that track keystrokes, speech patterns, facial expressions, or workflow decisions allow firms to convert formerly tacit or highly contextual knowledge into codified data, effectively reducing the knowledge advantage that workers hold. As the cost of knowledge extraction falls, firms may prefer this route, particularly in settings where workers have limited bargaining power.

The process of closing or eliminating the tacit knowledge gap is closely tied to the de-skilling of labor. When firms are able to capture and embed the specific skills required to perform a job into technology or standardized processes, they no longer need to hire workers who personally possess those skills. For example, when Starbucks replaced traditional espresso machines with automated ones, employees no longer needed to know how to grind beans, tamp espresso, or judge extraction by sight and taste. Instead, making a drink largely became a matter of pressing a button. Historically, reductions in the skill and judgment required to perform a job have been associated with lower wages and reduced bargaining power for workers, as documented in (Autor; Thompson on expertise).

There are clear historical precedents for this pattern. As discussed above, early twentieth-century Taylorist management practices sought to systematize and standardize skilled labor by closely observing workers, breaking tasks into measurable components, and shifting control over work processes from labor to management. Likewise, historical accounts from plantation economies describe how overseers used contests among enslaved workers to discover the limits of human output, effectively using competition to extract workers' knowledge.⁵⁰ In both cases, firms increased their power over workers by narrowing or eliminating workers' informational and skill advantages.

c. C) Failure to adopt beneficial AI systems

A large worker-firm knowledge gap can produce not only degraded AI or exploitative AI, but could also lead to delayed or foregone adoption of AI systems that could otherwise be broadly useful. If workers believe that AI adoption might threaten their economic well-being, they may rationally resist adoption, even in cases where AI technology could raise productivity, reduce

⁵⁰ Caitlin Rosenthal, *Accounting for Slavery* 96 (2018).

errors, or improve quality.

Such resistance can take the form of strikes, legal challenges, or political advocacy that slows or blocks implementation. For example, in 2024 the International Longshoremen's Association (ILA) organized a strike in which thousands of dockworkers walked off the job, demanding a ban on autonomous and semi-autonomous port equipment such as cranes.⁵¹ If fully implemented, such a ban would likely reduce port productivity, as there is little dispute that machines can lift and move cargo more efficiently than human workers.⁵² The eventual contract reflected a negotiated compromise: ports were permitted to adopt "semi-automation," in which machines perform core tasks under human supervision, coupled with guarantees that new human jobs would be created alongside each new machine.

Another way in which worker hesitation may lead to bad AI outcomes is the demand for "human-in-the-loop" mandates that are motivated primarily by employment protection, but where evidence is lacking that this provides improved productivity or safety. In trucking, for instance, labor groups have pushed for rules that require a human driver to remain in autonomous vehicles.⁵³ Similarly, in medicine, professional groups such as the American College of Physicians have emphasized that clinical decision-making AI models should remain supportive, and should always require human oversight.⁵⁴ Although these mandates are often framed in terms of safety and trust, in some cases they may also be mechanisms for limiting the extent to which AI replaces professional judgment with codified, firm-owned decision rules.

⁵¹ Andrea Hsu & Scott Horsley, *Dockworkers go on strike, snarling traffic at East and Gulf Coast ports*, NPR (Oct. 1, 2024), <https://www.npr.org/2024/10/01/nx-s1-5133391/dockworkers-strike-east-gulf-coast-ports-shipping>.

⁵² The stated efficiency argument against automation was that it left work more vulnerable to cyberattacks.

⁵³ *California bill to have human drivers ride in autonomous trucks is vetoed by governor*, AP (Sep. 23, 2023), <https://apnews.com/article/california-self-driving-trucks-newsom-veto-b585941fa05fc2c6b3a1965bceec3ee73>.

⁵⁴ *ACP Recommends AI Tech Should Augment Physician Decision-Making, Not Replace It*, Am. Coll. Physicians (June 4, 2024), <https://www.acponline.org/acp-newsroom/acp-recommends-ai-tech-should-augment-physician-decision-making-not-replace-it>. While augmentation may be appropriate in many clinical settings, there is also research showing that it may be better (for patient outcomes) to grant AI models autonomy over certain types of cases, rather than allowing humans to overrule AI recommendations. See, e.g., Nikhil Agarwal et al., *Combining Human Expertise with Artificial Intelligence: Experimental Evidence from Radiology 2-6* (NBER Working Paper 31422, Nov. 2025), <https://doi.org/10.3386/w31422>.

These concerns are not new. Historically, worker resistance has often delayed or shaped the adoption of technologies that were widely regarded as productivity-enhancing. In the early twentieth century, for example, automatic, push-button elevators were technically reliable well before they became widespread, yet elevator attendants remained common for decades. Historical accounts emphasize that continued human operation reflected employment concerns and social expectations about safety and service, rather than doubts about the underlying technology.⁵⁵ More recently, the adoption of automated teller machines (ATMs) faced opposition from bank workers' unions concerned about job losses, even though ATMs ultimately reduced costs, expanded banking access, and increased consumer convenience.⁵⁶ In each case, resistance did not stem from uncertainty about technological effectiveness, but from concerns about how the gains from automation would be distributed between firms and workers.

To understand the best legal intervention to balance the competing concerns of workers, firms, and social value from automation, we need to get more granular about how different workers are situated with respect to automation of their cognitive work. Part II lays out in greater detail how different workers are currently situated in the landscape of automation risk.

II. LABOR AS DATA AND ITS ROLE IN AI DEVELOPMENT

The effects of surveillance-enabled AI will not be uniform across workplaces. Different types of work vary in how easily labor data can be collected, in the value of labor data for building good AI models, and in how substitutable labor data from different individuals is. Part II introduces a framework for thinking systematically about these differences. It identifies the key dimensions that shape how surveillance-enabled AI may impact workers and firms and explores how each dimension gives rise to distinct policy and governance challenges.

Part A lays out a stylized taxonomy of worker types based on characteristics

⁵⁵ Bernard, Andreas (2014). *Lifted: A Cultural History of the Elevator*. Gray, Lee (2002). "From Ascending Rooms to Express Elevators." *Technology and Culture*, 43(3): 570–596.

⁵⁶ Bessen, James (2015). "Toil and Technology." *Finance & Development*.

relevant to how automatable such workers' work may become. This taxonomy lays out distinct individual starting positions, for workers. Part B then explores how these individual worker types may take advantage of (or overcome the shortcomings of) their individual positions, by considering how these positions will affect worker bargaining in the shadow of automation.

A. Taxonomy

The impact of surveillance-enabled AI depends critically on individual and collective worker agency—the degree to which workers can influence how data about their labor are collected, used, and monetized. Worker agency depends on two dimensions: *replicability* and *substitutability*.

1. Replicability

The first dimension, *replicability*, captures how easily a firm can build an AI model of a worker's labor. Replicability depends on two factors: the inherent legibility of a worker's processes and their ability to resist surveillance.

Some forms of work are highly legible meaning that its key inputs and outputs can be readily captured by monitoring technologies. Call-center agents and drivers fall into this category: recordings of customer interactions or detailed telemetry on steering and braking provide rich data for constructing a digital model of their work. By contrast, other kinds of work are far less legible. A mathematician's reasoning, for example, takes place largely as unobservable cognitive processes: notes or proofs offer only partial glimpses of the underlying thought, making it very difficult for an employer to capture the data necessary to reproduce it without that worker's consent and participation.

Beyond legibility, replicability also depends on how much discretion workers have to resist or shape surveillance. Call-center agents, whose tasks occur entirely on employer-controlled platforms, have little ability to prevent data capture. Truck drivers have somewhat more autonomy and, as Karen Levy documents, some have sought to disable or evade on-board monitoring devices.⁵⁷ Many knowledge workers retain greater discretion still: consultants, for instance, can decide whether to store client notes on employer servers or on personal devices. Differences in both legibility and resistance thus jointly determine how readily firms can translate lived labor into digital replicas. Workers who can limit or shape what is recorded retain more agency.

⁵⁷ DATA DRIVEN (2023)

2. Substitutability

The second dimension of worker agency, *substitutability*, captures how many other people could reasonably perform the same job. This varies with the importance of worker identity, heterogeneity in skill among workers, and the scarcity of their expertise and local knowledge. Worker substitutability is particularly low when the identity of an employee matters for the quality of the output. A movie studio cannot easily replace Tom Cruise in a Mission Impossible sequel, or a real estate developer might value a building designed by Renzo Piano, not merely for the design itself, but also for the prestige associated with the designer.

Substitutability may also be low when identity doesn't matter, but when quality is variable, and top talent is rare. For example, knowledge data generated by a top-performing salesperson or software engineer may produce far more value than an average one, even though both nominally perform the same job. When a worker's knowledge data is unique, that worker will tend to have more agency determining how her data will be used and whether or how she will be compensated.

Finally, worker substitutability can be low when a worker has exclusive access to local knowledge, such as sole understanding of client demands or firm-specific work processes that require extensive on-the-job training. The value of local knowledge can degrade quickly in some cases, increasing worker substitutability. For example, in cases where consultants have high client churn. However, in certain cases, the environment creates enduring local knowledge capture, such as trade-secrets a firm must protect with a limited number of trusted employees or in the context of non-compete laws that limit hiring another worker with similar domain expertise. Workers whose labor data are more substitutable will in turn have less individual bargaining power, and, as explored in greater detail below, must band together collectively to leverage the power of their knowledge data.

3. The Four Worker Types

Combining these dimensions yields four broad worker types.

a. Templates

At one extreme are *templates*: workers who are both relatively easy to replace and easy to replicate. Most customer service representatives, especially those handling routine inquiries, fall into this group. Individual workers are often

randomly assigned to customer queries, and there is considerable turnover, making it important that workers be relatively interchangeable.⁵⁸ Their tasks are also highly legible: nearly all customer interactions occur through monitored channels that are automatically recorded, leaving little room for workers to prevent or shape surveillance. Similar dynamics apply to other forms of routinized, surveillable labor, such as software engineering, data entry, and warehouse logistics. Firms are likely to hold substantial bargaining power over individual template workers because they can capture knowledge data and construct AI models with minimal cooperation or consent.

b. Artisans

At the other extreme are *artisans*: workers who are both difficult to replace and difficult to replicate. This group includes top talent in knowledge-intensive professions—star scientists, consultants with deep client relationships, or surgeons whose names carry independent value. These workers are not easily substitutable because their contributions, whether real or perceived, are distinct from those of others in their field. They are also difficult to replicate because their performance often depends on tacit knowledge and other inputs that are hard to observe or codify. A consultant’s reasoning process, for example, cannot be reconstructed from her final slide deck, nor can her instinct for how a client will react to a proposal. And if she wished to resist datafication, she would have numerous opportunities to withhold or obscure the expertise that makes her valuable. As a result, artisan workers are more likely to retain some agency over their labor data.

Between these extremes lies two additional groups of workers: *idols*, who are difficult to replace but easy to replicate, and *vessels* who are easy to replace but difficult to replicate.

c. Idols

Idols include star actors and influencers, workers who are unique because of their fame or identity, but who nonetheless may be easy to replicate because their output is easy to digitize into training data. For example, the filming of a single movie produces enough visual and audio data to create a digital replica of an actor that could, in principle, be reused indefinitely for other roles.

⁵⁸ See Oliver E. Williamson, *Transaction-Cost Economics: The Governance of Contractual Relations*, 22 J.L. Econ. 233, 254-57 (1979), <https://www.jstor.org/stable/725118>; Oliver E. Williamson, *The Economic Institutions of Capitalism: Firms, Markets, Relational Contracting* (1985).

d. Vessels

Finally, many “average” workers, such as routine consultants, teachers, and retail employees are vessels. While these workers perform similar work and often hold few unique skills relative to others in similar roles, they are difficult to replicate because their effectiveness depends on judgment or intuition that is difficult to capture in workplace recordings. We refer to this group as vessels because they hold tacit routines, heuristics, and relationships that firms rely on but cannot easily extract or codify.

*

Two important things are worth noting explicitly about these categories. The first is these categories apply to workers at specific points in their careers, and can change both across workers and over time. Absent any further industry regulation, an extra in a film just beginning her career is most likely a template worker. A movie star, however, is an idol who would likely be able to negotiate some terms of employment.⁵⁹ If an actor is cast in a break-out role, then his status in our taxonomy would change from template to idol. Similarly, one can imagine junior associate lawyers beginning their careers as vessels and eventually becoming artisans.

Importantly, the status of workers can also change due to changes in technology that make them more replicable. Most surgeons, for instance, are currently vessels because the tacit dimensions of their work are difficult to capture on video and because surgical robots are only starting to be developed, making their work difficult to replicate. Yet there are already companies using videos of procedures to train newer robotic surgery arms.⁶⁰ If this happens, then surgeons may begin to look more like templates, akin to how large language models made customer service work much more replicable by machines.⁶¹

B. *Datafication Effects Across the Taxonomy*

In this section, we examine the consequences of datafication for workers in each of the quadrants described above absent some legal or other intervention. To be clear, these consequences cannot support drawing conclusions about the overall welfare effects of worker knowledge datafication. Widespread

⁵⁹ For example, Harrison Ford was de-aged using AI technology for a recent Indiana Jones movie, and was presumably paid for it.

⁶⁰ See, e.g., Ji Woong Kim et al., *Surgical Robot Transformer: Imitation Learning for Surgical Tasks*, <https://surgical-robot-transformer.github.io/> (last visited Mar. 3, 2026).

⁶¹ CTK

datafication of worker knowledge may impact individual tasks but need not translate directly into its impact on jobs as a whole. For example, if datafication enables firms to codify the knowledge required to perform a worker's core tasks, one possibility is that the worker is displaced. Another possibility, however, is that the worker reallocates time toward tasks that are more complex, interesting, or better compensated. Because there is little evidence regarding how jobs may be reorganized in response to the development of more AI models, it is difficult to predict how workers may be impacted more broadly.

Beyond its effects on workers, datafication may also affect consumer welfare. These effects could be positive, for example by improving quality or reducing costs and enabling lower prices. Alternatively, datafication could lead to lower-quality products or services through an over-reliance on cheap but degraded AI models.

For these reasons, in what follows we organize our analysis around differences in workers' core tasks, using the quadrant framework above to clarify how datafication is likely to affect workers with distinct combinations of substitutability and replicability.

1. Templates

The most successful workplace uses of generative AI so far have come for tasks where human workers have generated large digital records of their labor: customer service and coding. The resulting models work well because these jobs are highly legible and generate large numbers of similar examples completed by many different people. In other words, workers in these settings are often templates.

Firms hold far greater bargaining power over template workers because they can capture labor data and construct AI models with minimal cooperation or consent. For example, movie actors have reported being asked to provide body scans with little understanding of the ownership rights related to these biometric data.⁶² These workers report having limited ability to push back: "Actors are, by and large, people pleasers. To have a standoff about scanning when you are

⁶² See Michael Savage, *supra* note [x]. In another creepier example, employees at X uploaded biometric data to train their Xai chatbot.

<https://www.theguardian.com/film/2025/oct/17/ai-data-scanning-film-tv-actors-crew>

another creepier example about employees at X training Xai chatbot that we could use instead: <https://www.theverge.com/news/814168/xai-grok-ani-employee-biometric-data>

in the midst of a scene annihilates your creativity, engenders fear that you will never work again, that your agent will drop you. So you comply.” These practices occur in other sectors as well: according to one industry estimate, “70% of new employee contracts” are expected within the next few years to include provisions granting employers broad rights over workers’ digital personas.

In these cases, workers are not offered any additional compensation for the use of their labor data. Yet even under existing law, workers could, in principle, demand higher pay if they knew their data would be used to train AI systems. In practice, however, template workers are unlikely to receive meaningful compensation. Because they are close substitutes for one another, the marginal value of any single worker’s data is low (and in many cases, effectively zero) and templates have not yet built up scarce local knowledge on-the-job. If one worker refuses to have their data used, the firm can simply draw from another.

Knowledge data for individual template workers thus becomes a commodity: interchangeable, and cheap. This dynamic is especially clear in jobs where producing training data *is* the job itself. Indeed, companies already outsource massive quantities of data labor to low-wage countries. In the Philippines, for example, workers label images, annotate videos, or correct text at pay rates below local living wages. The generation of commodity labor data now also extends to physical labor, such as on “hand farms” where workers record themselves performing repetitive labor like folding towels in order to generate training data for robotic AI.

Even when workers’ knowledge data are not identical—say, because some employees are more skilled or because their tasks differ, such as handling different types of customer complaints—existing legal and market structures still push compensation downward. This is because even data from average or lower-skilled workers can improve an AI system. For instance, transcripts from routine customer service calls still help chatbots recognize common problems and refine their tone.

If a lower-skilled worker agrees to share their data for a modest payment, the firm can use that information to make its model more capable overall. Once the model improves, the firm’s need for better data declines. A top performer who holds out for more compensation may find that the firm now has a “good enough” model trained on cheaper data instead.

This dynamic creates a negative *competition externality*: when one worker sells their data, they reduce the bargaining power of everyone else. Each worker’s

individual choice is individually optimal, but does not take into account the harm it does to other workers in the same position.

The consequences of this dynamic are that firms are likely to be able to turn the expertise of template workers into AI capital at little extra cost. Depending on how these AI models are used, this could lead to substantial reductions in the employment required to perform the same tasks (if the models replace workers entirely) or reductions in the skills required of future workers performing such tasks (if models support novice workers in performing work that more experienced workers once did).⁶³

2. Artisans

At the other extreme of our taxonomy, artisan workers are likely to experience datafication in a different way. In labor markets, workers with rare, high-value expertise typically command high wages and are capacity-constrained: a top consultant can only handle a limited number of cases at once, and a leading physician can see only so many patients in a day. Precisely because these workers are both expensive and scarce, firms have strong incentives to datafy their work. A digital system that could capture a star consultant's reasoning or an expert clinician's diagnostic process would lead to productivity gains, both from reducing the cost of accessing that expertise and by scaling it across more customers.

Yet, firms are very limited in their ability to datafy artisans because their work is difficult to replicate without cooperation. An artisan who wishes to resist datafication has a wide range of practical tools at their disposal.

First, most workplace surveillance tools simply cannot capture the heart of what these workers do. For example, a screen recording of a senior consultant's laptop would reveal typed notes or spreadsheet edits, but not the tacit reasoning behind them: why a diagnostic question was asked at that moment, how the consultant read the room in a high-stakes meeting, which client relationships were being leveraged to unblock a process, or how they integrated small cues from conversation into a larger strategic judgment. These are the kinds of expert

⁶³ We note that deskilling or replacing tasks is not the same as deskilling or eliminating jobs: jobs are bundles of tasks, and the impact of AI will depend on how jobs are reorganized. For example, technological change typically reshuffles which tasks are done by people versus machines rather than wiping out whole occupations at once. The task-based literature shows that automation substitutes for routinized tasks while complementing other, often nonroutine, activities; how pay and employment move then depends on how the job is reorganized after this re-bundling.

moves that generate real value, and they do not show up on a keystroke log or webcam feed.

Second, even in domains where digital traces exist, artisan workers can often withhold or fragment the relevant data in ways that materially reduce its usefulness. A physician can sketch differential-diagnosis notes on paper before entering the formal chart; a consultant can draft sensitive analyses on a personal device before transferring only the polished conclusions; a researcher can brainstorm on a whiteboard that is never recorded. These methods do not reduce performance, but they do prevent the employer from systematically capturing the full logic of the worker's process.

In contrast, consider what might happen if artisans were motivated to create digital models of their expertise. When artisans see value in digitizing their skills, they are more likely to provide model builders with richer data that captures their actual reasoning. They might share annotated versions of draft work, explain why they made particular decisions, or supply contrasting examples that show both effective and ineffective approaches to a task. With this level of cooperation, the model-development process can begin to resemble *reinforcement learning with human feedback* (RLHF), a common method for refining AI systems. In an RLHF setup, the model produces an output and the artisan responds by indicating what is correct, what is missing, and how they would handle the same problem. Repeating this feedback loop allows the system to learn judgment that would be illegible given routine surveillance data.

Because of the gains to cooperation, artisans are likely to retain substantial agency over their labor data, relative to templates. For example, an experienced consultant, surgeon, or litigator might negotiate higher wages, royalties, or separate contracts in exchange for providing detailed feedback, curated examples, and structured reasoning used to build an internal “expert system.” Unlike the case with templates, artisans are also less likely to compete with each other in providing data because they are not perfect substitutes for one another; each brings a distinct style, client base, local knowledge or domain specialization that firms may want to capture. Indeed, there are currently a variety of services that enable content creators (such as management thinkers, social media influencers, or public educators) to digitize aspects of their expertise.

3. Idols

The situation with idols is at first similar to that of artisans. Because idols are highly unique—consider one star actor versus another—it is likely that idols

will be able to negotiate rights for the use of their digital replicas. Indeed, there are already cases in which actors and influencers have been compensated for the use of their digital likenesses or data, such as agreements governing digital deaging in film or the licensing of AI-powered chatbots trained on an influencer’s persona. In some ways, the potential benefits of datafication may be larger for idols precisely because their work is so readily replicable by AI tools. An actor’s digital replica, for example, could participate in the production of many more films, advertisements, or localized versions of content than the actor could physically perform. If governed by clear contractual protections, such arrangements could allow idols to scale their labor, earn income detached from time constraints, and extend the commercial life of their creative output while retaining control over how their persona is used.

In practice, however, career dynamics make this outcome difficult to achieve. Few idols begin their careers as idols: most famous actors start as unknown performers in small roles—that is, as template workers. As discussed above, template workers typically have little bargaining power, making them unlikely to negotiate substantial compensation or strong usage restrictions for their digital data. Yet because digital data are durable, early-career agreements can persist long after a worker’s market value has changed. As a result, even if a template worker later becomes an idol, they may already have licensed the use of their data at a “template price.” Absent meaningful limits on how digital replicas can be reused, an actor would need to negotiate—at the outset of their career—compensation commensurate with the full expected value of their future performances. In practice, only established stars possess the bargaining power to demand such terms. Consequently, tomorrow’s idols are unlikely to secure these protections early on, meaning that by the time they achieve stardom, they may have already relinquished control over valuable aspects of their digital likeness.

4. Vessels

As with artisans, it is difficult for firms to independently replicate the labor of vessel workers. Consider a schoolteacher: a classroom camera can record lessons, but this captures only a thin slice of what makes the teacher effective. It does not encode what she chooses not to say, how she adapts instruction to particular students, or the reasoning behind pedagogical choices. Training a genuinely capable model would therefore require the teacher’s active cooperation in providing feedback on student misunderstandings, annotations explaining instructional intent, or counterfactual explanations of alternative

approaches. This suggests that even where vessel workers are, in principle, replaceable, they may retain some bargaining leverage because firms depend on their cooperation to generate usable labor data.

However, vessels seeking to be compensated for their individual labor data face competition both from other vessels, as well as from more distinguished artisan workers in the same field. For example, an early career management consultant may attempt to negotiate compensation for the use of her project documentation or decision processes, but her employer may opt to hire a different consultant who is willing to supply similar data without additional compensation, especially if her client-specific knowledge or on-the-job training is of limited value. Alternatively, the firm may prefer to pay a premium for data from a senior partner whose work product is perceived as higher quality or whose reputation makes their work processes more valuable.

These dynamics suggest that datafication may generate limited returns for many vessel workers while increasing within-occupation inequality: once work is codified and reusable, modest differences in perceived quality or reputation can be scaled across many clients and projects, making it rational to pay substantially more for even marginally better performers. This is the logic of “superstar” markets, often invoked to explain rising inequality in settings such as CEO compensation.⁶⁴ Taken together, these forces imply that while vessel workers may retain some leverage in the datafication process, the economic gains from datafication are likely to accrue disproportionately to a small subset of top performers rather than being broadly shared across the occupation.

**

What Part II illuminates is that the risks and benefits of AI adoption are not evenly distributed across the labor market. Workers face diverse threats and opportunities from AI. Some workers face pressure from firms investing in “templating” their work—turning non-template forms of expertise into forms of labor more readily exploitable for automation (and associated deskilling). For example, surgeons being studied to inform robotic surgical arms. Others face significant intra-sector competition, as the promise of “winner take all” automation pits workers against one another to capture the outsize gains that come from being automated and scaled. This risks producing negative competition externalities, as discussed above. Excessive worker competition to

⁶⁴ Rosen 1981; Gabaix and Landier 2008

capture all the gains of automation can exacerbate still another problem: the expertise pipeline that makes skilled, valuable workers from less skilled ones over time. Pipeline threats are particularly acute where workers begin as one type (e.g., templates) and require both firm resources and personal experience to build up the expertise needed to evolve into different worker types (e.g. artisans or vessels) at which point some workers may produce new forms of valuable and innovative labor. Cutting off that pipeline puts future forms of expertise and workplace innovation at risk.

These dynamics pose diverse challenges to developing the best possible automation, too. Template workers may have the least leverage to affect automation quality by withholding labor data—but even among well templated work, eliminating workers’ ability to withhold *any* tacit skill is rare. In the aggregate, even small incentives for increased worker cooperation may meaningfully improve models over time, or reduce the costs in building good models. Intense intra-worker competition resulting in winner-take-all dynamics may favor the bold over the best: allowing early movers to lock in benefits but not necessarily rewarding the highest quality workers, especially when fierce competition to automate chokes off time in which quality differences may begin to appear. This risk becomes especially acute over the medium term: the best worker of today may be automated at the expense of the systems needed to produce, let alone automate, the best worker of tomorrow. Especially in settings where it is difficult (or impossible) to know in advance which template workers will evolve into say, the best idols or vessels, automation risks letting labor markets ossify around existing automated work and underinvest in future labor innovation.

III. KNOWLEDGE GUILDS

Collective data governance mechanisms are a promising way to solve the misalignment issues that generate risks of worker exploitation, degraded AI, and foregone beneficial AI canvassed above.⁶⁵ By securing workers’ rights to benefit from their automated expertise, at a minimum knowledge guilds can create buy-in to encourage workers to contribute to AI development and act as mechanisms

⁶⁵ Cullen, Li Data supra note [x]. Viljoen has argued a related claim: that as a legal matter, collective data governance mechanisms are the only way to secure data subjects’ legal interests in social data. Viljoen supra note [x].

to distribute and share the productivity gains of knowledge capital. If correctly designed and implemented, knowledge guilds thus make possible credible firm commitments to share the benefits of automation. This reduces the risk of worker exploitation and, relatedly, undercuts the motivation for workers to block socially beneficial AI.

Prominent AI leaders agree with us, at least behind closed doors. In unsealed documents from *Bartz v. Anthropic*, Dario Amodei, the co-founder and CEO of Anthropic, wrote a document arguing that the creators behind AI's training data should get paid.⁶⁶ He said it was a “real and important concern” that big companies were training on human labor to “extract concentrated profits.” These models would outcompete the humans behind the training data in the long run. It would therefore be fair, even “better for everyone” to share profits with them. Others have similarly agreed that collective bargaining and value sharing mechanisms are needed to for workers to maintain the value of their common knowledge resources.⁶⁷

More ambitiously, knowledge guilds can produce a more diverse ecosystem of institutions contributing to digital innovation and experimenting with how to share the benefits of digital innovation more widely. In the Conclusion we survey the diverse implications that knowledge guilds may have across different industries. Despite some variation across worker types, we believe the introduction of robust and widely available knowledge guilds will be broadly beneficial across worker types. But first, this Part lays out what knowledge guilds are, what rights they secure, and a range of options for introducing them into workplaces.

A. *Digital Sovereignty and Digital Association*

The knowledge guild concept has considerable flexibility in terms of form and content. But to prevent the suboptimal AI outcomes discussed in Part IB above, knowledge guilds must meet two essential conditions. First, knowledge

⁶⁶ *Bartz v. Anthropic*

⁶⁷ Nicholas Vincent, Matthew Prewitt, Hanlin Li, *Collective Bargaining in the Information Economy Can Address AI-Driven Power Concentration*. <https://arxiv.org/pdf/2506.10272>

guilds must recognize the principle of workers’ digital sovereignty. Second, they must recognize that workers’ digital sovereignty depends on other workers and thus must be secured collectively.

1. Digital Sovereignty: three equal data interests

Digital sovereignty encompasses three distinct interests that people have in their data: privacy interests, control interests, and financial interests.⁶⁸

Privacy. First, privacy refers to the interest people have in limiting the who, what, and how of information about themselves being shared. Legal privacy covers a variety of specific claims and is operationalized across legal domains via varying mechanisms⁶⁹. To be clear, privacy rights in their current form are almost always too weak to serve their social function.⁷⁰ Too often privacy is a legal fiction. A single fleeting indicia of consent sanctions broad and extensive data uses.⁷¹ The failure to exercise a difficult-to-access opt out licenses broad,

⁶⁸ The use of ‘interest’ as opposed to ‘right’ is intentional. As explored in greater detail below, knowledge guilds secure workers’ interests in data not via individual rights to data, but via individual associational rights to form knowledge guilds. That being said, privacy rights enshrine current legal rights—so where appropriate, the terms “right” and “rightsholders” are used to refer to such privacy rights.

⁶⁹ Most data collected about people is governed by contractual terms of service, subject to the Federal Trade Commission Act’s Section 5 and state consumer-protection oversight. *See* 15 U.S.C. § 45(a)(1) (2018) (prohibiting “unfair or deceptive acts or practices in or affecting commerce”). All states have incorporated similar consumer-protection clauses into their civil codes, and state attorney general offices use their enforcement authority under such statutes and myriad other state privacy laws to regulate consumer digital terms and services.

A series of sector-specific privacy laws have granted additional rights to consumers over particular kinds of data. These include the Health Insurance Portability and Accountability Act of 1996 (HIPAA), Pub. L. No. 104-191, 110 Stat. 1936 (codified as amended in scattered sections of 18, 26, 29, and 42 U.S.C. (2018)); the Children’s Online Privacy Protection Act of 1998 (COPPA), 15 U.S.C. § 6501-06 (2018); the Family Educational Rights and Privacy Act of 1974 (FERPA), 20 U.S.C. § 1232g (2018); the Fair Credit Reporting Act of 1970 (FCRA), 15 U.S.C. § 1681 (2018); alongside a few prominent state laws like Illinois’s Biometric Information Privacy Act (BIPA) 740 ILL. COMP. STAT. 14/10, 14/15 (2021); and the California Consumer Privacy Act of 2018 (CCPA), CAL. CIV. CODE §§ 1798.100-1798.199.100 (West 2021).

⁷⁰ Cite to a few privacy critiques.

⁷¹ On consent and the legal theory of legitimacy for click-through or standardized consumer contracts, see, for ex- ample, NANCY KIM, *CONSENTABILITY: CONSENT AND ITS LIMITS* 91-116 (2019); NANCY KIM, *WRAP CONTRACTS* 126-46 (2013); and MARGARET RADIN, *BOILERPLATE: THE FINE PRINT, VANISHING RIGHTS, AND THE RULE OF LAW* *passim* (2012). On click-through digital consent more specifically, see Elizabeth Edenberg & Meg Leta Jones, *Analyzing the Legal Roots and Moral Core of Digital Consent*, 21 *NEW MEDIA & SOC’Y* 1804, 1804-05 (2019); Woodrow Hart- zog & Neil

ongoing data collection.⁷² Privacy rights are often bundled with access to key services, such that the only way to exercise one’s privacy right is to forego the service writ large.⁷³

Feeble as they are in practice, in theory privacy rights enable rightsholders to exercise a “stop” function: they empower data subjects to cut off a flow of information, or to refuse to provide information at all. In this sense, privacy rights index the negative claims individuals have *against* informational uses, and to draw the boundaries that limit which information may be collected from them.⁷⁴

In the U.S., data rights tend to be synonymous with privacy rights.⁷⁵ But this approach is incomplete. Data subjects do not only have privacy rights *against* compelled intrusion or expropriation of data. Data subjects also have affirmative interests in deciding how data about them will be used and to enjoy the fruits of its value. These affirmative rights to data control and to data value are just as important as privacy rights against informational intrusion. Simply put, data subjects deserve the capacity to say ‘no’ to data collection and use, but they also deserve the power to say ‘yes’—to employ their data in the service of their

Richards, *Privacy’s Trust Gap: A Review*, 126 YALE L.J. 1180, 1197-98 (2017) [here- inafter Hartzog & Richards, *Privacy’s Trust Gap*] (reviewing FINN BRUNTON & HELEN NISSENBAUM, *OBFUSCATION: A USER’S GUIDE FOR PRIVACY AND PROTEST* (2015)); Neil Richards & Woodrow Hartzog, *Taking Trust Seriously in Privacy Law*, 19 STAN. TECH. L. REV. 431, 434 (2016) [hereinafter Richards & Hartzog, *Taking Trust Seriously*]; Andrea M. Matwyshyn, *Technoconsent(t)sus*, 85 WASH. U. L. REV. 529 *passim* (2007); Elettra Bietti, *Consent as a Free Pass: Platform Power and the Limits of the Informational Turn*, 40 PACE L. REV. 307, 329-31 (2020); and Daniel J. Solove, *Privacy Self-Management and the Consent Dilemma*, 126 HARV. L. REV. 1880, 1883-85 (2013).

⁷² Elements of the CPRA restrict this, but it is common elsewhere. See e.g. software licenses where data subjects must agree to terms of service to use the product, including privacy terms, websites that will not load if you have ad blockers on, etc.

⁷³ Elements of the CCPA restricts bundling services with data collection, but again, it is common elsewhere.

⁷⁴ Several privacy theories offer more robust and complete theories of informational governance that are not merely negative. See e.g. contextual integrity, Priscilla Regan. Such theories aggregate what we here call privacy and control interests under the umbrella term privacy. We disaggregate them for two reasons. First, despite the intellectual attractiveness of these privacy theories, many privacy laws do not in fact protect this more robust conception of privacy. Second, privacy and control interests may at times conflict. Disaggregation helps surface the tradeoffs data subjects must make as they participate in digital life.

⁷⁵ Cite to Margot Kaminski on data protection and privacy in the US.

priorities and values. These affirmative interests are explained in greater detail below.

Control. Second, control refers to the affirmative interests people have in directing how information about them is used.⁷⁶ Control interests allow data subjects to exercise a “go” function: they empower data subjects to manage how data flows from them into the world. Data control captures the interest data subjects have in both the conditions under which their data is used, and the ends to which their data is put. Some privacy laws (and theories of privacy) are broader than simply rights against re-identification. Such accounts of privacy include elements that grant control or quasi-control rights.⁷⁷ However, data control is worth distinguishing because privacy (what we here call privacy) and control answer for distinct interests in data.

For instance, data may be formally private—deidentified or not used in a way that would implicate an individual qua herself—but may still violate a data subject’s control interests if such data was being used for purposes the data subject finds unethical or amoral. Just as one might decline to work for an employer who behaves illegally or unethically or resign from an organization that has engaged in actions one does not want to be associated with, individuals and communities may wish to withhold their data from certain parties or purposes for any number of legitimate reasons even if their privacy is otherwise

⁷⁶ The most intuitive legal means of securing control is of course to recognize a property right. There is a longstanding debate among scholars about whether the law ought to recognize a formal property right in data. One of us have written critically about endowing data subjects with *individual* and *full* ownership rights over their data. See e.g. Relational Theory, Data Property? However, others have written in favor of recognizing such rights. See James Grimmelman piece. Elsewhere, Viljoen has suggested that a form of data ownership that is collective in its structure and with limited alienability would be in line with what we propose here. See *Daten als Eigentum? in Umkämpftes Eigentum* (2025 eds. Niklas Angebauer, Jacob Blumenfeld and Tilo Wesche). By disaggregating the ‘bundle’ of property interests and emphasizing control (and exploitation), this proposal hopes to capture what is attractive and useful about property in this context, while discarding that which is not.

⁷⁷ Indeed, one of the earliest US theories of privacy was of privacy as a form of informational control. See Alan Westin. However, in practice many privacy laws, if they are indeed enshrining a right of control, enshrine a very limited, binomial control right: the right to say yes or no to full access or no access at all. What this grants the data subject is a veto right. The control interest we envision here is much richer and affirmative interest: the interest of the legislator in crafting and negating the shape of legislation and in setting the legislative agenda, not the interest of the executive to exert a one-time veto right over what comes across the desk.

guaranteed.

Importantly, control interests may be violated by inaction as much as action. For instance, consider a knowledge guild that partners with a research agency to conduct medical research using health data and no research is undertaken. This does not violate data subjects' privacy, but it does thwart their affirmative intentions to put their data towards a desired use.

Payment. Third, financial or payment interests refer to data subjects claim to their share of value their data helps produce. As discussed at length in Parts I and II above, the digital economy—and AI in particular—depends on information that workers produce and use to train and improve the systems. Yet under most current legal arrangements, workers (and consumers, whose data also improves systems) lack legal rights to claim and negotiate for a 'dividend' or share of the wealth these systems generate.⁷⁸

Digital sovereignty is a general concept as it applies to consumer data and worker data alike. However, this Article is particularly concerned with a particular kind of workplace data: namely the data that captures the expertise and knowledge a worker brings to bear on her tasks. Thus, the aim of knowledge guilds is to secure a specific kind of data sovereignty.

2. Digital Association

Digital sovereignty articulates the interests that data subjects (and by extension, data workers) have in a digital society. But a crucial challenge remains: *how* can these interests be secured and exercised?

The nature of information markets means that exercising individual data subject rights, even comparatively strong ones (like those Europeans enjoy), does not exert meaningful discipline on the platforms, AI companies, and employers who use worker data to improve their AI products. Given how information markets work and how data produces value, granting workers individual rights of privacy, control, and payment will not protect the data sovereignty interests at stake in workers' knowledge data.⁷⁹ Instead, these

⁷⁸ Cite to data dividend act.

⁷⁹ Some privacy interests are indexed via individual privacy rights, and rightfully so. We see

interests are best secured by granting workers associational rights to form and join knowledge guilds.

There are two kinds of arguments that support associational protections for knowledge data. The first are arguments grounded in how informational markets work generally. These arguments imply that associational or collective forms of data governance are preferable for all kinds of data, not just knowledge data. Beyond the context of worker knowledge, digital sovereignty is best secured via collective protection and expression for several reasons.

First, the value of most digital data is relational and combinatorial.⁸⁰ Data become economically and analytically useful only when aggregated with data from many other individuals; their predictive and productive power increases with scale and diversity. Up to a point, therefore, data exhibits increasing returns to aggregation. At the same time, this relational structure implies that any single individual data stream typically has very low marginal value within a large corpus. The withdrawal of one person's data rarely meaningfully reduces the performance or commercial value of a large training dataset. Moreover, because many attributes can be statistically inferred from existing population-level data, the informational footprint of any one individual is often partially reconstructible even in their absence. These structural features limit the effectiveness of purely individual exit as a mechanism of digital sovereignty and support the case for collective data protections generally.

Relatedly, the control and privacy interests of data subjects are highly interrelated. Prominent privacy theorists have long recognized privacy as a good best secured via broad social protections, due to the contextual nature of informational flows and the privacy externalities that flow from data's relational qualities.⁸¹ Priscilla M. Regan emphasizes the common stakes of privacy, given market forces that make it difficult for any one individual to have privacy unless

such privacy rights as complementary to, not substitutes of, the privacy interests workers have in their labor data. The main point here is the inverse of the same lesson: privacy rights are important, but they cannot protect and stand in for the other interests in data that workers have.

⁸⁰ Viljoen, *Relational Theory* (611). Posner and Weyl *Radical Markets*

⁸¹ Nissenbaum, *Contextual Integrity*, on privacy externalities, see Katherine Strandburg *Free Fall: The Online Market's Consumer Preference Disconnect*, 2013 U. CHI. LEGAL F. 95, 143; Daniel J. Solove, *Privacy Self-Management and the Consent Dilemma*, 126 HARV. L. REV. 1880, 1889-93

a minimum is guaranteed to everyone.⁸² Helen Nissenbaum develops an account that encompasses both privacy and control as appropriate information flow, where what constitutes appropriate (or inappropriate) information sharing is determined by reference to socially developed norms in a relevant context, such as the workplace.⁸³ On this account, privacy is governed not individually but by the “web of constraints” that make up the various norms of social life.⁸⁴ Pauline Kim and Rachel Leavitt argue that collective data governance needs are *especially* strong for workers.⁸⁵ In their account, workers have a collectively shared joint interest—not merely individual interests that point in the same direction—in how employers use their data to manage workplace conditions.⁸⁶

These arguments support associational protections for data subjects generally. But workers and the production of knowledge data also generate specific challenges that make knowledge guilds attractive. In Part IIIB below, we build the case for knowledge guilds as a specific form of collective data governance by considering how they respond to the diverse pressures different workers face due to replicability and substitutability. Knowledge guilds can be usefully distinguished as a category of data association from other forms of data collectives along many dimensions. One worth noting here is the nature of their intervention in the political economy of AI innovation. Kim and Leavitt, for example, make general claims about workplace data practices, that might involve generalized surveillance practices applied to a workplace. Our argument below focuses particularly on the role knowledge data plays in turning labor into (AI) capital.

In the workplace setting, the primary benefit and justification for knowledge guilds is to help workers collectively secure their financial interests in data value. However, workers also have privacy and control interests in their knowledge data. By securing the full suite of worker data interests, knowledge guilds can recognize the range of motivations and concerns that may either prevent—or if protected, empower—workers to contribute their knowledge to automation.⁸⁷

⁸² Priscilla Regan at 225.

⁸³ Nissenbaum, Contextual Integrity.

⁸⁴ *Id.* at 128.

⁸⁵ Pauline Kim and Rachel Leavitt, Data Rights for Data Workers, 106 Boston University L. Rev. (2026) at 6, 42.

⁸⁶ *Id.*

⁸⁷ Elizabeth Anderson makes a related argument regarding workplace governance more generally. She notes that workers valued unions not only for higher wages, but also as a means

3. The Basic Features of a Knowledge Guild

Knowledge guilds protect worker's data sovereignty interests in knowledge data via a special form of work-specific data association. Below, we explore in greater detail different motivations for forming knowledge guilds that will have distinct benefits and drawbacks for differently situated workers.

Here we sketch out features that a well-functioning knowledge guild ought to have. Knowledge guilds share several essential features with traditional craft guilds, adapted for informationalized skill. As such, knowledge guilds should include the following components.

Seamless pooling mechanism. First and foremost, a knowledge guild must either pool together knowledge data directly or collectively manage and enforce the legal rights its members enjoy over their knowledge data.

Civic data trusts such as those managed by the Open Data Institute offer a real-world example of how data pooling can work.⁸⁸ Civic data trusts pool community-generated data and collectively govern such data, including settings terms of access that platforms must honor to gain access. As such, civic trusts can bargain for data governance when individual members lack market power. The principle is the same for knowledge guilds, but the stakes are more pressing. Workers are not contributing demographic preferences or location data, but hard-won expertise on which their livelihoods depend.

For workers that do work where the knowledge data is useful independent from context, their data may be hosted by a knowledge guild directly, and then only accessed by firms on defined terms. For example, consider a surgeon who specializes in a particular kind of spinal surgery, and whose knowledge data a firm wants to use to train a robotic surgery arm. The surgeon's knowledge data may need to be adapted for the robotic arm to use, but their knowledge data's usefulness does not depend mingling it with firm information, it depends almost entirely on the skill and experience of the surgeon. For firms to then gain access to this knowledge data, they need to agree to the terms of access set and

to exert greater control over the conditions of the workplace. Private Government (2017).

⁸⁸ Cite to Open Data Institute

enforced by the knowledge guild. Such terms will reflect the surgeon and the guild's negotiated bargain with such firm on this data's value but also set terms on data control and privacy. A 'technical' knowledge guild like this can work especially well to allow the surgeon, or other workers whose knowledge data value is independent from firms, to scale their expertise across firms. This kind of access model is not unlike how firms gain access to AI models now. They pay for access to the foundational model technology and then adapt it for their firm's context.⁸⁹

But for many kinds of work, local firm knowledge and worker knowledge must be used together to train a model. In these settings, it will make less sense for guilds to act as a technical intermediary hosting and managing access to knowledge data, than for guilds to act as a legal intermediary. Here, knowledge guilds can negotiate and enforce the legal terms of how knowledge data will be generated, managed, and used to train models. Such knowledge guild terms can provide for knowledge data what labor union contracts provide for labor: setting floors on how value is shared between workers and firms and negotiating the control and privacy conditions that shape how knowledge data contributes to model development. As a general matter legal guild protections apply to workplace settings where guild-covered knowledge data is being used in a guild-covered manner, regardless of whether the knowledge data itself sits in a guild server or at a firm.

Some guilds may combine the technical and legal guild options into a hybrid guild system. For example, one could imagine SAG-AFTRA providing a technical guild option to manage access to likeness and voice data of actors for future works. Future projects would then lease access to actors' knowledge data subject to the terms set by the knowledge guild and the actor. This may be attractive for actors to, for example, lend their voices to animated projects or narrate videos for causes they care about without concern about such data's abuse, while they then free up time to physically work on other projects. But SAG-AFTRA could also offer actors, and the estates of deceased actors, legal knowledge data protections for actor knowledge data. This could include knowledge data already encoded in past projects. Clear legal protections for

⁸⁹ Indeed, over the medium term one can imagine technical guilds using their collective informational resources to train their own models that the guilds can then themselves provide to firms.

knowledge data would counteract claims by studios that may allege a right to use existing records of past projects to produce “digital scabs” from such actors’ past performances. Legal protections would also ensure that voice and likeness data generated from future physical collaborations is protected by guild terms.

Tracking knowledge contributions. In both settings above, knowledge guilds will need to either directly deploy or require the means of tracking worker knowledge contributions. Enforcing guild protections over knowledge data depends fundamentally on a guild’s ability to access the information necessary to quantify the knowledge contributions of its members. Traditional craft guilds collected records of contracts, wages, and working conditions across a trade to assess and enforce the going terms of a marketplace. Similarly, knowledge guilds must be capable of tracking the stream of inputs its members provide to AI models—and crucially, of measuring the value those inputs generate.

Developing systematic and high-quality measurements of knowledge contributions is essential for guilds to negotiating over sharing their value. The most tractable unit of measurement is hours of labor saved, which map naturally onto wages. Hours of labor saved would provide a legible basis for assessing the contribution of any individual or cohort of workers. Over the longer term, knowledge guilds and firms will have additionally developed a set of market terms for different knowledge data and kinds of usage from which to bargain over future contributions.

Contribution tracking is related to the knowledge data pooling mechanisms above. Technical guilds, where access to knowledge data is mediated via a guild interface, can build in this kind of data tagging as a part of its access protocols. Legal guilds will also require this kind of information, but as described below, can require data tagging and ensure firm compliance via auditing or monitoring rights.

And as a practical matter, in neither case does value accounting depend on building out entirely new infrastructure. Cooperation with the foundational models already deployed in many workplaces could facilitate accounting in many settings, by allowing knowledge guild members to upload their knowledge data as a distinct, tagged training stream. This would preserve the provenance of member contributions and make it possible to attribute downstream model improvements to their source.

With measurement comes negotiating power. Once a knowledge guild can credibly demonstrate—in hours saved, error rates reduced, or productivity gains realized—the aggregate value that member contributions have generated, it is in a stronger position to negotiate the price of that knowledge data's value on behalf of its members. Armed with this evidence, a knowledge guild can begin to function as guilds historically have: settings floors for acceptable terms, enforcing standards of contribution and compensation across an industry, and ensuring that the gains from automation are shared with those whose knowledge made those gains possible in the first place.

Enforcement mechanism. Knowledge guilds need to have teeth to secure protections. There are two pathways to endowing knowledge guilds with bargaining authority: enabling legislation that explicitly grants workers bargaining and association rights over knowledge data, or via voluntary options to create knowledge guilds via existing legal mechanisms.

The strongest pathway is through enabling legislation that would 1) recognize workers' rights to join a knowledge guild, 2) specify the legal forms guilds may take in relation to pre-existing membership organizations or as stand-alone legal entities, and 3) require (or encourage via tax benefits or liability exemptions) firms to bargain with knowledge guilds over the terms of accessing, producing, and using knowledge data.

Legislation publicly recognizing knowledge guild rights could take a variety of forms. It could impose greater liability for firm data misuse and then offer direct bargaining with knowledge guilds as a means to obtain exemption from liability, similar to how platforms enjoy copyright liability immunity under the Digital Millennium Copyright Act (DMCA) if they create a DMCA-compliant takedown system.⁹⁰ Knowledge guilds as a means of liability immunity would be especially attractive should legislative attempts to overturn Section 230 of the Communications Decency Act prove successful.⁹¹

Guild-enabling legislation could specify membership thresholds that

⁹⁰ DMCA cite.

⁹¹ CTK on various legislative efforts to overturn S230, or otherwise limit its broad immunity.

determine how many members a guild needs to have before firms are required to recognize and negotiate with them. Legislation may also create exemptions for firms based on size, or better, for firms whose work involves minimal AI adoption. Enabling legislation could, but need not, specify whether guilds obtain at the enterprise (e.g. “Amazon engineers”) or sectoral (e.g. “online influencers”) level. In fact, we think leaving this unspecified may foster a more robust marketplace of guild options. Once guilds form and attract members, the market for guild membership itself can reveal what organizational form suits worker and firm purposes best.

However, in our view, enabling legislation should ensure that knowledge guilds remain most responsive to their members, and make clear that fiduciary duties to members supersede any fiduciary duties that may arise should a knowledge guild have a corporate form that allows third party investment. Thus, enabling legislation should ideally impose strong conflict-of-interest standards that ban or cap the ability for owners with a stake in a firm employing knowledge data to also own a stake in a knowledge guild. So, for example, Amazon or its subsidiaries should not be able to claim a stake in the “Amazon engineers” knowledge guild (or be subject to a strict ownership cap), and third-party for-profit stakes in knowledge guilds should in general be subject to increased scrutiny and monitoring for conflicts of interest.

In settings where workers enjoy relatively strong bargaining power and can overcome coordinating challenges, workers may privately come together to form knowledge guilds even in the absence of enabling legislation. Guild-like options exist that would allow some workers to secure some protections now.

First, groups of workers can form *private data trusts* to pool knowledge data and collectively manage terms for accessing that knowledge data resource. As noted above, civic data trusts and other forms of data collectives already exist and can serve as a helpful model for workers looking to pool together and manage access to their data together. Civic data trusts, such as the Open Mobility Foundation, offer a model of a third-party trust. In this hub-and-spoke model, the data trust sits “outside” any one organization.⁹² The trust creates an API to collect data that flows to and “sits” with the data trust, and member organizations can get insights and analysis back. Other trusts, common in the

⁹² Open Mobility Foundation

health context, operate slightly differently. These operate more like an intermediary layer or platform. For example, many health organizations upload health data into a trust, which then manages access by members or other parties for research use by tagging and managing data permissions. Both model may be adapted for knowledge data depending on context.

Existing labor groups are best positioned to use data trusts to create more protections for workers now. For example, SAG-AFTRA could stand up a data trust as part of its offerings to members, in line with the examples described above, and rely largely on its existing bargaining frameworks to develop differentiated payment rules for different tiers of members' knowledge data. Professional licensing organizations like the AMA could extend its licensing requirements to include compliance with stated knowledge guild conditions for professionals' knowledge data use. Mission oriented employers like faith organizations and NGOs may voluntarily set up third party data trusts to manage workers' knowledge data. Independent data trusts that have separate fiduciary obligations to members' data can act as a form of estate planning, should the organization go under or be acquired by a future owner with a different sense of mission, and act to credibly constrain how that organization's workers' knowledge data may be used in future partnerships.

However, the private data trust option will be limited in practice by several conditions. First, direct data pooling will only benefit workers that do work whose value is fairly firm-agnostic, as noted above. Second, in the absence of a clear legal bargaining right, workers themselves will need to enjoy at least some bargaining power, meaning they are in a field either with limited competition, or high levels of solidarity. This is in some tension with the first condition, since the more firm-agnostic knowledge data is, the more workers may compete against one another to capture gains. Finally, workers will either need to spin up a data trust themselves, or partner with groups who have expertise in setting up and managing data trusts, and commit to funding its management on an ongoing basis.

Another ready option for some guild-like protection is for worker organizations to disseminate *form contract language* for employment contracts that impose guild-like terms over knowledge data. The goal of such form language is to disseminate default norms around knowledge data value and to encourage widespread adoption. Worker associations or advocacy groups could develop

worker-friendly contract terms that workers can then negotiate for directly with employers. Ideally, such terms will make clear how the scope of employment will apply to their knowledge data and impose guild-determined data tagging protocols as discussed above. Such terms should further specify how workers' labor data may and may not be used, and what uses would trigger additional payment or value sharing for the worker and the guild.

The biggest drawback to form contract language is that it is entirely voluntary and individual and thus will largely track existing bargaining power between individual workers and firms. But in some settings form contract terms provide a powerful baseline for negotiations. For example, the National Venture Capital Association (NVCA) provides “model legal documents” that serve as an industry-wide baseline for tech startup investments.⁹³ By coordinating to provide sample documents from which such deals proceed, the NVCA proposed a set of default norms about the kinds of contracts considered typical for such deals, and to provide an easy mechanism for those norms to proliferate in the sector.

One step further is to create voluntary knowledge guilds that attract membership via member support services that include data use conditions alongside other forms of support, including data value monitoring services on behalf of members. Where private data trusts offer a private pathway to a technical knowledge data guild, this offers a private pathway to something more like a legal knowledge data guild.

For example, the Authors Guild (which provides a variety of services to authors, who to join the Guild must pay a membership fee), offers both contract review and model contract language for authors negotiating with publishers.⁹⁴ In fact, the Authors Guild already suggests language regarding use of works for training AI. The Authors Guild also suggests other terms that provide authors ongoing rights to a share of profits in (and in some cases, creative control over) derivative works. Such terms typically cover rights for an author should a creative work be adapted to other mediums (for example, adaptation to film or television). But the basic principle of an ongoing authorial right to derivations or adaptations of original work can be extended to the

⁹³ CTK

⁹⁴ Cite to authors guild.

‘derivative work’ of knowledge data into an automated system meant to supplement or replace a work task.

The extension of this principle seems particularly plausible in creative industries where (some) workers winning residual or ongoing claims to value is an established contracting norm. In both film and music industries, for example, artists have long negotiated ongoing contractual rights to payment from their films or music. Actors can negotiate for residual payments for their films and television. Musicians register their (copyrighted) work with the Mechanical Licensing Collective to ensure smooth royalty payments from digital music providers, and many partner with performing rights organizations (PROs) to license their performance rights and music publishers who handle licensing and registration, as well as offering many other services. Musicians usually obtain royalties through these contractual relationships, as the “songwriter’s share” from their publishing agreements.

Again, the biggest barrier to achieving such protections via voluntary membership organizations is that unless membership is sufficiently high, firms may have little incentive to bargain with guild members. Thus, voluntary mechanisms will face enduring challenges to achieving the structural market disciplining functions needed.

Internal representation, adjudication, and differentiation mechanisms. Finally, knowledge guilds need mechanisms to manage the different priorities of its members. Different workers will have different priorities regarding how to protect privacy, generate and price data value, and exercise knowledge guild control over how data is used. Knowledge guilds will need internal voting and representation rules to for determining guild positions, managing conflicts, and ensuring member accountability. We think predetermining the content of such rules in the abstract is ill advised. Guilds will determine over time the bylaws and governing articles that make sense for their members—who in turn will challenge these rules in ways that allow doctrines of guild governance to develop over time.

B. Implications: replicability and substitutability

As noted above, knowledge guilds combine two features relevant to protecting worker claims to labor data. First, data sovereignty names the worker

interests at stake in automation. To have teeth, and to be both justiciable and legitimate, knowledge guilds need to *entitle* workers to their data interests. Second, data association captures the group or collective nature of data interests that are thickly bound up with those of other workers. To impose market discipline on firm behavior, knowledge guilds also need to offer workers a *collective action mechanism* for pooling and expressing their data interests, as well as internal mechanisms for resolving tensions that may arise among workers when their interests conflict.

In our view, both entitlements and collective action are essential components of meaningful worker empowerment, but they protect different dimensions of labor. Workers will value each differently depending on two features of their work: replicability and substitutability. As discussed in Part II, replicability refers to how easily a firm can build an AI model of a worker's labor. This depends on how legible the worker's processes are in digital form and on how much control the worker has over the knowledge data they supply to their employer. Substitutability, meanwhile, concerns how easily a firm can obtain equivalent knowledge data from other workers.

Replicability primarily determines the importance of entitlements. When it is technologically feasible for firms to replicate workers' expertise without their consent, preserving meaningful worker power requires giving workers legally protected claims to their knowledge data. Substitutability, instead, shapes the need for collective action. When individual workers' knowledge data are highly substitutable, their individual bargaining power will be weak. In such settings, protecting workers' data sovereignty is unlikely to succeed through individual negotiation; it requires coordination across workers. In short, entitlements address technological vulnerability, while collective action addresses bargaining vulnerability.

1. Replicability and the value of entitlements

A formal legal entitlement to data would counteract worker replicability. Granting workers a legal right to control their data and to share in its value—either as a form of data labor, or as a new class of intangible property owned by the worker—would, at least in theory, insulate (or insure) workers against their

replicability.⁹⁵ Entitlement to labor data relocates the source of a worker's power to withhold data. Currently, how much withholding power a worker enjoys depends on the technical (and investment) limits of a firm's ability to surveil workers and to render more work legible. This source of withholding power is not under a worker's control and is vulnerable to future erosion. A legal entitlement to data relocates the source of a worker's power to withhold from the technical (in)capabilities of the firm to the law.

It is thus unsurprising that proposals for data entitlement rights are perennially common not just for workplace data, but for data generally.⁹⁶ One of the most enduring forms data ownership proposals take is to recognize data as a new class of intangible property, by formally recognizing transferable, fungible, property rights for data subjects in their personal data.

Propertarian entitlement reforms are typically put forward in one of two forms. Most propose a property right in data that entitles the data subject to sell access to her data streams (like a license or easement) or transfer full ownership rights in a given data exchange. For example, Sir Tim Berners-Lee who developed the core protocols of the World Wide Web, later founded Solid Project, to give data subjects "true data ownership," which he saw as central to its aim to "radically change the way Web applications work today."⁹⁷ Similarly, Former U.S. presidential candidate Andrew Yang proposed a right to data property in his campaign platform⁹⁸ and later launched the Data Dividend Project to pressure companies such as Facebook and Google to pay users a "data dividend" for the wealth generated by their data.⁹⁹

A variant on the property claim envisions data-generation as a form of labor operating in a free-labor market. In distinction with our argument here, this proposal views *all* data production—not just the knowledge data produced while performing labor—as labor.¹⁰⁰ In this variant, data collection entitles the data

⁹⁵ In prior work, Viljoen has written about granting property entitlements to data. See *Relational Data, Phenomenal World*, discussing how different forms of entitlement mechanism would be more or less vulnerable to exploitation by technology firms.

⁹⁶ Viljoen, *Phenomenal World*, *Data as Property?*

⁹⁷ What is Solid?, Solid Project at MIT, <https://solid.mit.edu/>

⁹⁸ See *Data as a Property Right*, Yang2020, <https://www.yang2020.com/policies/data-property-right/>

⁹⁹ Data Dividend Project, <https://datadividendproject.com/aboutus>.

¹⁰⁰ Data generated about labor should not be confused with the proposal that all data

subject to command a wage in a data-labor market. Jaron Lanier was among the earliest to propose data-as-labor.¹⁰¹ His work influenced Glen Weyl (an economist) and Eric Posner (a legal scholar) whose project Radical Markets includes a proposal to create a labor market for data.¹⁰² In their view, casting data as labor is preferable over data as property, because labor as a general category of input to value creation better captures how individuals contribute to the data economy. Entitling users to command a wage for their data-labor is thus necessary to restore (or rather, create) a functioning market for user contributions.¹⁰³

Both variants of this proposal—data as labor or data as property—work by “capitalizing” a data-asset—coding it via law with protections and features that recognize data’s role in contemporary wealth creation.¹⁰⁴ This transforms data about the subject into an asset that can generate wealth *for* the subject.¹⁰⁵ It thus allows data producers to reclaim at least one interest in data: its value. For many proponents of formalizing legal property rights in data, this proposal is rooted

produced via data subjects should be considered labor. While terminology is shared, the arguments are distinct. Entitling all data producers with a wage right treats all data production as labor; our claim is that data generated *about* labor—in particular data that mediates worker knowledge and expertise—is a source of worker-produced value that warrants greater worker authority and control over that value.

¹⁰¹ Cite to Lanier; 2014

¹⁰² Eric A. Posner & E. Glen Weyl, *Radical Markets: Uprooting Capitalism and Democracy for a Just Society*, Princeton University Press (2018), 208-09.

¹⁰³ Eric A. Posner & E. Glen Weyl, *Radical Markets: Uprooting Capitalism and Democracy for a Just Society*, Princeton University Press (2018), p 209; Imanol Arrieta Ibarra, Leonard Goff, Diego Jiménez Hernández, Jaron Lanier & E. Glen Weyl, *Should We Treat Data as Labor? Let’s Open Up the Discussion*, Brookings Inst. (Feb. 21, 2018), <https://www.brookings.edu/blog/techtank/2018/02/21/should-we-treat-data-as-labor-lets-open-up-the-discussion/>

¹⁰⁴ As Katharina Pistor aptly describes, once data is conceived of as an asset of any kind in law, the conceptual distinction between capital (K) and labor (L) is reduced. Katharina Pistor, *The Code of Capital*, Princeton University Press, (2019), p 11. In law, both concepts treat data as the subject of an exchange relation between data subject and data processor for data’s alienable value. As a legal conceptual matter, alienable value in the form of L is easily turned into K with a bit of legal engineering. Take for example partners in a limited liability partnership (LLP). They contribute their labor to the corporate entity as in-kind services and take out dividends as a shareholder in lieu of a salary, thus benefitting from the better legal protections and a lower tax rate afforded K for the same exact work that would be performed where it coded as L instead. *Id.* at 48. Beyond law, the concept of “human capital” in corporate finance, organizational sociology, and labor economics also collapse the conceptual distinction between the contributions of capital assets and labor power to production. See e.g. Claudia Goldin, *Human Capital*, in *Handbook of Cliometrics* (2016); Samuel Bowles and Herbert Gintis, *The Problem with Human Capital Theory—a Marxian Critique*, 65 *American Economic Review* (1975) pp74-75.

¹⁰⁵ Posner & Weyl, *supra* at 205-07.

in pragmatism: data is already being “coded” as quasi-capital in law, through a combination of contractual agreements and trade secrecy law.¹⁰⁶ The problem is that these informal propertizing legal mechanisms enable data holders to hoard the wealth data produces, excluding the individuals from whom data originated.¹⁰⁷ On this account, the issue is that data producers lack formal property rights that entitle them to benefit from their data assets.

Legal entitlements to worker data would affect different worker types differently. As an initial matter, the more replicable a worker is, the more important a legal entitlement to labor data would become for that worker. As discussed above, in the absence of a legal claim to withhold data, such workers have little power to withhold data otherwise—by definition, firms can monitor their activity and create a model to capture their skills to a reasonable degree of success without their cooperation. Thus, templates and idols would (in theory) stand to gain the most protection against expropriation from legal entitlements to withhold data than would vessels and artisans. For instance, call center workers are fully recorded, so they have little technical scope to withhold data. On the other hand, consultants can easily withhold their reasoning. Granting legal entitlements to withhold data would place the call center worker on (somewhat) more equal footing with the consultant, by granting legal recognition to the call center worker’s control and value interests in his data. Importantly, legal entitlements alone cannot address the larger market incentives to make work legible.

As will be discussed in detail below, data entitlements absent a collective enforcement mechanism will provide little value to template workers. But such entitlements can, as an initial matter, formalize a legal means of withholding data that replicable workers currently do not have.¹⁰⁸ In the case of idols, who possess more bargaining power, entitlements to withhold data may have considerable value on their own. Indeed, some actors are already negotiating for a share of the value from the automation of their likeness—a form of negotiation that formal entitlement will likely only accelerate and distribute among idol worker

¹⁰⁶ Amy K Law of Informational Capitalism; Julie E. Cohen.

¹⁰⁷ Julie E. Cohen, *Between Truth and Power*, Oxford University Press (2019), p 63; For a detailed treatment of how assets are coded in law to become capital, see Katharina Pistor, *The Code of Capital* (2019), pp 2-3.

¹⁰⁸ Of course in some instances, some worker data is arguably covered by copyright and other intellectual property regimes. But a broad data entitlement would cover many types of worker generated data beyond the scope of intellectual property.

types more broadly. Data sovereignty rights seem especially important to clear up cases where the power to negotiate an idol worker's data rights is currently ambiguous. For instance, when a famous actor has passed away, their estate and any studios that own works in which they performed may have competing views about how (or whether) to use the deceased actor's likeness for future projects.

On the other hand, firms stand to gain more from legal entitlements for vessels and artisans. Replicability tracks the productivity returns to the worker stake in data.¹⁰⁹ These workers are hard to replicate without their active cooperation. So, they may not be as concerned about automation (and labor data expropriation) and thus place less value on a legal right to enshrine the power to withhold data that they already have. Thus, workers with low replicability—like vessels and artisans—need legal entitlements less to protect against coerced labor data expropriation.

By the same token, these workers may be those most likely to supply more knowledge as a result of legal entitlement, resulting in more productive automation. Vessels and artisans currently lack systematic incentives to spend time and effort training AI models rather than engaging in their normal labor. Granting low replicability workers entitlements is one important means of facilitating and incentivizing them to cooperatively share more labor data. Explicit legal rights to the benefits of that cooperation can help firms make credible commitments to such workers that they will be compensated for their cooperation. Thus, settings with low replicability workers may be where firms and workers are most aligned in benefitting from worker data sovereignty. Firms benefit from the extra data effort workers might put in to training models, and workers benefit from being able to get a share of the productivity gains generated by better AI.

Most of the existing proposals for data entitlement envision granting data producers individual entitlement rights without also tying such entitlement rights to legal mechanisms to pool and exercise them collectively. In our view, entitlement is an important element of protecting worker's knowledge sovereignty, but unlikely to broadly achieve the goal of worker knowledge sovereignty alone. A legal entitlement alone is not a silver bullet. It does not

¹⁰⁹ In other words, the less replicable a worker is, the more her active cooperation is needed to produce automation with the highest possible productivity gains

erase or fully stand in for background market conditions in a labor market. These background conditions will thus significantly affect the scope and power of the entitlement.

As noted above, in some settings, firms and workers may both benefit from individual data entitlements. In such settings, individual entitlements even in the absence of a collective pooling mechanism *can* meaningfully alter how firms and these workers share the productivity gains of AI. But there are several caveats that limit how widely this bargaining position applies in the labor market. First, even for artisans, in settings where workers are price takers and face strong labor market competition the value of a legal entitlement to withhold data will go down as workers compete away its value. If workers feel there is a large gap between the value of their knowledge data contribution and its compensation, they again may systematically under share labor data. Relatedly, in the absence of strong incentives to the contrary, firms have an incentive to invest in making more types of work more replicable. Over time, this shrinks the types of work where, by its nature, workers enjoy more bargaining power over automation productivity gains.¹¹⁰ Thus, workers granted individual entitlements to withhold their labor data may not be able to enforce market-wide changes in either current firm behavior or worker behavior via individual entitlements alone. To produce market-wide changes, workers need collective action mechanisms to resist competing away value and to reduce the incentive of firms to invest in innovation that relies on worker expropriation for productivity gains.

2. Substitutability and the value of collective action mechanisms

Pooling data interests via a bargaining unit (what we call a knowledge guild), would smooth differences across substitutability. Forming knowledge guilds would, at least in theory, insulate (or insure) workers against their substitutability. Knowledge guilds relocate the source of workers' power to negotiate for a share of value their labor data generates from (highly variable) features of their particular labor market to the law.

Currently, a worker's power to negotiate their share of automation value

¹¹⁰ Indeed, as discussed in Part [X], this is the long sweep of history of scientific management.

depends on several features of their work. As noted in Part II above, substitutability depends on how important a worker's identity is to their work (as is often the case in creative settings like acting or design), variation in skill across workers (as is common in work that requires extensive training), and the scarcity of expertise (which may reflect both marketwise scarcity, but also importantly includes work where hyper-localized knowledge is important).

Unlike the case of replicability, where the *source* of the power to withhold data largely reflected features of work outside a worker's control, workers' power against being easily substituted is at least partially a function of their own efforts. Take variation in quality, for instance. Part of what distinguishes a star basketball player from a more ordinary one is innate talent differences—but part of it is also about effort. But like replicability, workers' power against substitution is also vulnerable to future erosion, as automation threatens the talent pipeline for certain kinds of low substitutable work and pushes more workers into high substitutability work over time. A ready legal mechanism to pool worker data value thus similarly relocates the source of workers' bargaining power—from labor market conditions that may empower some workers but imperil many others, to collective bargaining mechanisms anchored in law.

Again, it is thus unsurprising that many scholars have argued that—even beyond the case of labor—data sovereignty is best secured collectively. Several data property theories cast data as a collective, even public form of property.¹¹¹ While knowledge guilds for labor data do not require theories of collective or jointly held worker data entitlements, they are compatible with such theories. Moreover, as discussed above in Part A, there are compelling reasons to think even if workers have individual interests in data sovereignty, strong collective action mechanisms are needed to secure and express those interests.

Collectivizing data value management via knowledge guilds would affect different workers differently. When workers are largely performing the same functions, they lose more from unregulated competition between themselves to supply labor data. Knowledge guilds are thus highly beneficial for template workers, since the collective mechanism of the guild enables template workers

¹¹¹ Roberta Fischli on collective property; Open Data Institute; Data commons (commonwealth). Indeed, Viljoen suggests that certainly in many instances, data interests obtain at the population, not the individual level. Recall Kim and Leavitt make this case explicitly for worker data.

to obtain value from their data entitlements. In the absence of a guild, template workers, such as call center workers, that attempted to withhold data value would simply be replaced by workers who do not.

Data guilds are also highly beneficial for many vessel workers. For these workers—who are not easily replicable but are highly substitutable, guilds also help prevent such workers from competing away any gains from labor data. While firms do need vessels' active cooperation in supplying labor data for automation, any individual vessel worker that demands too much value for their labor data may be substituted for a worker willing to supply labor data for a lower price. Data guilds may be especially and unambiguously valuable for workers where nearly all the workers in that field are vessels.

Different sources of substitutability can cut across industries and affect the value of knowledge guilds in important ways. Sometimes a vessel and an artisan do the same type of work; their differences in type reflect variations in their skill levels.¹¹² Where the degree of substitutability stems from variance in skill, then there is a risk that for example the “star Supreme Court litigator” (i.e. the artisan) will capture all the gains of automation at the expense of the more-substitutable average litigation associate (i.e. the vessel). In these settings, the risk is less that fellow average litigators will compete away data value, and more that a “star” worker will capture all the gains of automation. In these settings, it is worth thinking carefully about how to balance spreading the gains of automation more widely among workers with ensuring that uniquely skilled/talented workers are still appropriately incentivized to contribute their expertise (and thus get the best-possible automation).

Other variations in substitutability are less vulnerable to intra-worker competition.¹¹³ In fields where the source of workers' low substitutability stems from the importance of their identity, or their highly specialized and scarce expertise, the value of workers' data is not as zero sum. For example, the value of a model of Tom Hanks does not meaningfully reduce the value of a model of Christian Bale. Both models may be widely deployed in future films but nevertheless describe distinctive goods. Other workers' low substitutability

¹¹² This describes fields where there is vertical differentiation—the same type of work done with varying degrees of skill. Examples include consulting, teaching, litigating and sewing.

¹¹³ This describes fields where either individual identity of the worker is important, or where the nature of the work means that the worker's expertise is scarce.

stems from the unique blend of their expertise with that of a firm in a way that provides highly localized forms of value. For example, consider a highly specialized engineer who works at a firm to develop a novel medical device. This worker is not easily substitutable in part because of their specialized training, but primarily due to the symbiotic relationship between their expertise and the proprietary firm technology blending to produce value. A model built from this engineer's experiments perfecting this medical device does not meaningfully threaten the value of a model of another researcher's experiments at a different firm.

In both these settings, artisans and idols may face little incentive to join knowledge guilds, since they benefit less from collective protection against competition. On the other hand, because their data value cannot meaningfully replace that of other workers, these workers also face less incentive to capture the outsize gains available from defecting from collective bargaining. Particularly where worker data value is highly related to its ability to blend with firm assets, it is important that knowledge guilds don't place too many constraints on firm and worker informational blending.

3. Putting it all together

In sum, different workers will value the distinct features of knowledge guilds—entitlement and collective action—differently. Template workers have the most to gain from knowledge guilds in their strongest form. Highly replicable and substitutable, granting template workers an individual entitlement to their data alone is unlikely to empower them to meaningfully negotiate a greater share of automation value. Template workers would particularly benefit from mandatory (or widely adopted) knowledge guilds across the labor market, that minimize the ability for firms to pay individual template workers to defect. By the same token, template workers are unlikely to benefit from piecemeal legal protections and voluntary knowledge guilds, such as form contract language or voluntary data trusts. Template workers face considerable collective action challenges in forcing industry wide standards if such standards are not separately legally required, or heavily incentivized for the firms not to circumvent.

Idols benefit from entitlements but arguably need collective bargaining mechanisms to gain from such entitlements less. In theory, a star idol (e.g., Tom Cruise or a star architect) can individually negotiate for some value of their

knowledge data in the absence of a guild. If guilds impose ceilings on compensation, particularly famous idols may even think that guilds prevent them from negotiating for their true value. However, the collective nature of guilds can offer idols other attractions. For one, many workers may join guilds before they become idols; indeed, the pipeline for future idols may depend on a guild, since many start as templates. And guilds can offer many compensation tiers to reflect workers' different productivity contributions. Since idols are generally not substitutes for one another, they also don't have much incentive to 'defect' from the guild to out-compete possible rivals. Therefore, getting idols on board with knowledge guilds likely depends on allowing individual guild members some leeway to negotiate higher value. Thus, in creative settings, knowledge guilds can create a floor for template workers, while providing detailed "payment tiers" for workers as they make the progression from template to idol.

Vessels need entitlement less, though firms (and by extension vessels) stand to gain from entitlements. However, vessels will benefit from guilds as collective action mechanisms. In these settings, guilds can prevent negative competition externalities that force workers to compete away value. They can also soften the effects of "superstar" artisan workers capturing all the gains from automation at the expense of other workers. Vessel workers will thus benefit substantially from guilds. Because vessels enjoy the ability to withhold data, in fields where all or most workers are vessels, workers may be able to have reasonable success with voluntary guild arrangements, especially in fields with reasonably high levels of solidarity.

However, other fields will face tensions between vessel and "star" artisan workers. In these settings, knowledge guilds will need to structure compensation such that "star" workers are still incentivized to provide their expertise. Not allowing higher compensation for higher skill could risk under providing the best expertise to train "best possible" models.

Artisans arguably stand to gain the least from guilds, as they are not particularly replicable or substitutable. Indeed, as discussed above, in some cases guilds may benefit vessel workers at the expense of "star" artisan workers in the same field. On the other hand, firms stand to gain from artisans joining guilds, as it creates a credible mechanism to gain their cooperation in making their expertise more conducive to automation. Future artisans may similarly benefit;

many artisans start out their careers as templates or vessels (for instance, first year associate lawyers or consultants, or medical residents). Moreover, it may be very hard for a worker (or firms) to identify *ex ante* which workers will become a star or not, in which case guilds act as a form of insurance for workers. This is one class of uncertainty that guilds resolve. Other forms of uncertainty may arise when it is hard to locate or attribute expertise. In such cases it may appear that a single “star” worker is the source of all the value, but in reality their expertise relied on the cooperation and input of other more “average” workers around them. In all such cases, guilds can act as insurance against uncertainty about the source of current and future expertise. Nevertheless, in some cases it may be beneficial to allow artisans to “opt out” of standard guild terms, or to create tiered reward structures within guilds to ensure career pipelines still reward—and thus develop—future idols and artisans.

CONCLUSION

The total and real-time tracking of workers, and the use of this information to train AI models creates new vistas for firms to replace labor with capital. But real time worker data flows also create new affordances for more socially beneficial arrangements. As shown in prior work, workers, at least under some conditions, have an important role to play in determining the form and value of AI automation. In many settings, it would be more effective to train an AI system in collaboration with workers. The importance of worker expertise to automation can, and this Article argues, should, entitle workers to share in the productivity enhancing gains that automation promises. However, as detailed in Part II, current workplace conditions are not conducive to achieving those conditions.

Knowledge guilds might meaningfully address many of the uncertainties raised in Part II, to achieve fairer and/or more efficient forms of workplace automation. As detailed in Part III, knowledge guilds entitle workers to knowledge data sovereignty and enable them to collectively bargain for its protection. Guilds may take a variety of legal forms—such as a data trusts, legal worker organizations, an element of an existing member organization. We think they would be most widely adopted and broadly useful if explicitly enabled via legislation, but workers have some options in some contexts to form guilds even without legislative intervention.

Knowledge guilds are not a silver bullet. International labor or trade rules will become more important if a country succeeds at imposing domestic knowledge data standards. High domestic protections will be cold comfort if firms exploit international knowledge data instead. Additionally, knowledge guilds will not stand in fully for other labor protections. For instance, robust rights to knowledge data may not protect workers from other exploitative workplace data practices.

Nevertheless, knowledge guilds provide a response to the growing problem of how to facilitate AI innovation that is broadly beneficial. First, by reducing uncertainty and securing forms of value sharing, knowledge guilds can increase workers' willingness to share their knowledge to train models. Second, if firms must negotiate with guilds for knowledge data in a market, knowledge guilds can help to price the value of knowledge contributions to productivity across different settings. This would direct some AI investment towards knowledge guilds for obtaining and using such knowledge. Together, these effects can reduce the frequency and likelihood of suboptimal automation

Knowledge guilds can also help to smooth the bifurcation between high value and low value workers (the "substackification" of work), by partially subsidizing lower value workers. This becomes especially valuable as a form of insurance where workers may not know what type they are, and as a form of expertise pipeline conservation where low-value workers tend to be new entrants, and depend on being supported as they gain expertise.

Such smoothing may be particularly important as a redistribution mechanism if chronic problems with redistribution fail to share the relative gains of automation more widely. Knowledge guilds may also be important under conditions where firms (or the market generally) face structural uncertainty about worker quality over time. Finally, if knowledge guilds also function to facilitate the pipeline of talent discovery and have an incentive to signal promising new entrants, they can help prevent longer term underinvestment in innovation. Lastly, knowledge guilds can act as guards against 'bad bets' on automation getting crystallized by firms in settings where worker knowledge is chronically under- or overvalued or otherwise misunderstood.