AI and the Extended Workday: Productivity, Contracting Efficiency, and Distribution of Rents*

Wei Jiang, Junyoung Park, Rachel Xiao, Shen Zhang†

First draft: January 29, 2025

Current draft: August 15, 2025

Abstract

This study examines how occupational AI exposure affects employment at the intensive margin, i.e., workday length. Using individual-level time diary data from 2004–2023, we find that greater AI exposure, whether from the ChatGPT shock or broader AI developments, is associated with longer work hours and less leisure, particularly non-screen-based leisure. Guided by a principal-agent framework, we find evidence for three channels: higher marginal productivity from AI-human complementarity, improved contracting efficiency from AI-enabled monitoring, and lower worker reservation utility, reflected in declining job satisfaction unrelated to salaries or employment risk. The workday extension is most pronounced in competitive labor and product markets, where weak worker bargaining power shifts productivity gains to firms or consumers. These findings challenge the view that technology reduces labor burdens, showing that it can erode work-life balance.

Keywords: Artificial intelligence; time allocation; work-life balance; principal-agent model

JEL Classification: D22, E20, J01, J22, J24, L11, O3

*This study benefits from comments and suggestions from Bo Bian, Honglin Ren, Yihui Wang, Ting Xu, and the participants at the following seminars and conferences: Baruch, Fordham, GSU, PBC Tsinghua, SUFE, UBC, AI in Finance Conference (Maryland), CICF Annual Meeting, 2025 CUHK–Fudan JRCIF Conference, and UNSW Corporate Finance Workshop. All errors are our own.

†Jiang: Emory University Goizueta Business School, NBER, and ECGI, wei.jiang@emory.edu. Park: Auburn University Harbert College of Business, jzp0170@auburn.edu. Xiao: Fordham University Gabelli School of Business, rachelxiao@fordham.edu. Zhang: Fordham University Gabelli School of Business, shenzhang@fordham.edu.

1. Introduction

Artificial intelligence (AI) was initially conceived with the goal of making human work and life more interesting, fulfilling, and less laborious. Paired with other technological advances like automation, AI has the potential to boost productivity, enhance job satisfaction, and promote a healthier work-life balance. Nevertheless, empirical evidence regarding AI's impact on both work and leisure remains inconclusive. While much of the discussion has centered on AI's capacity to displace labor in some scenarios and generate new roles in others (e.g., Felten et al., 2019; Webb, 2019; Acemoglu et al., 2022; Kogan et al., 2023; Hampole et al., 2025), relatively little attention has been given to how AI reshapes work on the intensive margin—particularly its effects on work time, contracting efficiency, and the distribution of productivity gains. This paper aims to fill these gaps by analyzing the micro-level impacts of AI on time allocation, drawing on detailed individual-level time diaries collected from 2004 to 2023. Through this examination, we investigate how AI exposure influences work supply at the intensive margin and assess its broader implications for firm valuation and economic outcomes.

The relationship between occupation exposure to AI and work time is a priori ambiguous. For any given task, AI-driven automation and efficiency improvements should theoretically shorten task duration. Additionally, wealth creation boosted by technology should entice individuals to allocate more time from work to leisure, provided that leisure is a normal good. However, the classical principle-agent model (notably Holmstrom and Milgrom (1987)) provides a rich set of predictions in a setting where a worker optimally allocates his effort based on the production process, monitoring effectiveness and personal preferences. AI's impact on the potential to enhance productivity in diverse fields,[1] combined with its capacity to improve monitoring and productivity measurement, can result in heavier workloads and longer hours. This effect is expected to be more pronounced in competitive product markets,

[1]E.g., stock analysis (e.g., Gu et al., 2020; Lopez-Lira and Tang, 2023), legal practices (e.g., Casey and Niblett, 2016; Surden, 2019), music generation (e.g., Briot et al., 2017; Briot, 2021), and accounting (Commerford et al., 2022).

where businesses face escalating expectation from customers and pressures from competitors' enhanced capabilities; and in labor markets with relatively inelastic supply, where workers lack substantive bargaining power to adjust their schedules to their own advantage. Furthermore, AI's integration of real-time effort tracking and improved information availability frequently erodes the division between work and personal life, further contributing to extended working hours for some individuals.

Data from the American Time Use Survey (ATUS) provides a unique opportunity to test the hypotheses. The ATUS conducts a cross-sectional survey each year, with an average annual sample size of approximately 26,400 participants. Our sample spans two decades from 2004 to 2023. Respondents document their activities using detailed 24-hour diaries at 15-minute intervals, from which market-based work time, leisure time, and some special categories (such as education and child care) can be calculated, with reasonable variations for sensitivity checks (e.g., whether social activities at the workplace count as work or leisure). To attribute the changes in workday patterns to AI, we then measure each occupation's AI exposure based on the textual correlation between task descriptions and the content of AI-related patents using large language models. We further distinguish between complementarity and substitution relationships between AI and jobs.

The advent of ChatGPT toward the end of 2022 provides a natural experiment to test how workers change time allocation when their jobs are disrupted or complemented by the new AI technology. Workers in occupations with higher exposure to generative AI experienced a significant increase in work hours and a decrease in leisure time following the introduction of ChatGPT. An interquartile increase in AI exposure is associated with a 3.75-hour increase in weekly work time. This effect is particularly evident in occupations that are more complementary to generative AI and in regions where AI awareness is higher, as measured by Google search trends. Given that the general public was largely unprepared for the exact timing of ChatGPT and even more so for its advanced "human-like" capabilities,[2] the prolonged work-

[2]The surprise by the general public was evident from the comments on social media shortly after ChapGPT's launching. For instance, The *New York Times* article titled "A Conversation With Bing's Chatbot Left Me Deeply Unsettled," published on February 16, 2023, captured many of these reactions and reflected the broader

day is likely attributable to the new technology. Our finding that AI technology increases overall work hours challenges the common expectation that it would reduce work time by enabling workers to complete tasks more efficiently.

The same relationship holds over the full sample period for occupational exposure to general AI technology. After controlling for individual characteristics, alternative innovation exposure measures, and a comprehensive set of fixed effects—including industry × year, state × year, year-month, and day-of-week—an interquartile shift in occupational AI exposure is associated with an additional 2.25 hours of work per week in the cross-section. When further controlling for occupation fixed effects, the within-occupation effect remains sizable at approximately 50% of that magnitude. An employment analysis suggests that this extension of the workday is unlikely to be driven by task consolidation following workforce reductions.[3]

On the leisure side, reductions associated with AI exposure are primarily concentrated in non-screen-based activities such as relaxing, socializing, and traveling. This shift indicates that workers in AI-exposed occupations not only work more but also reallocate their leisure time toward screen-based activities, which are generally more passive and less restorative.

We test predictions from the Principal-Agent model along three dimensions: marginal productivity, monitoring efficiency, and reservation utility. The first set of tests shows that both work hours and wage rates increase monotonically with the level of AI exposure, in addition to net complementarity of AI. This pattern supports the hypothesis that AI-augmented productivity incentivizes workers to extend their working hours.

The second test examines AI's role in performance monitoring, using the rapid adoption of AI-driven surveillance during the 2020 pandemic-induced shift to remote work as a natural experiment. Among occupations that were ex ante feasible for remote work, defined by the absence of essential on-site requirements, workers in roles with higher exposure to AI monitoring technologies, particularly those aiming at direction and evaluation, experienced greater

public astonishment at the technology's capabilities.

[3]After our sample period, 2024–25 saw large-scale layoffs in the tech sector. Our default sample explicitly excludes this industry, and aggregate unemployment has remained relatively stable well into 2025. Thus, the inference that extended workdays are not primarily driven by fear of employment risk or task consolidation into a reduced workforce still holds.

increases in work hours following the shock. Crucially, this effect is not observed among the self-employed, for whom the principal-agent problem is moot. This serves as an effective placebo test, reinforcing the interpretation that AI monitoring contributes to extended work hours through AI-enhanced monitoring efficiency.

The final set of tests explores the extent to which the productivity surplus from occupational AI exposure has accrued to workers, and whether the cross-sectional variation in workers' "reservation utility" helps explain observed effort patterns. Although we confirm that firms benefit from the AI exposure of their workforce (evidenced by higher return on assets (ROA) and labor productivity), workers report lower satisfaction, particularly regarding work-life balance (based on Glassdoor reviews) even though AI exposure is associated with higher compensation. In other words, there is little evidence that productivity gains translate into improved worker welfare. In the cross section, we expect workers to capture a smaller share of the surplus (and work longer hours) when they have weaker bargaining power relative to their employers, or when their employers face limited pricing power vis-à-vis consumers.

Indeed, the extension of the work hours is more pronounced when workers have limited bargaining power-due to employer dominance in a regional and occupational hiring market, which restricts their information, mobility, and choices. Similarly, workers have difficulties extracting rents from technology-enabled productivity gains when the product market is competitive, leading to most of the rents being passed on to consumers and leaving little for firms to share with workers. In both scenarios, the reservation utility of workers (reflecting overall welfare in equilibrium) fails to keep pace with productivity gains during the AI boom, thereby undermining the income effect that would have otherwise induced more leisure and discouraged work.

Our study contributes to the rapidly growing literature that analyzes the impact of AI on the economy. A growing body of research (Autor, 2015; Felten et al., 2019; Webb, 2019; Acemoglu et al., 2022; Yang, 2022; Babina et al., 2024; Hampole et al., 2025) has uncovered various facets of AI's impact on businesses and employment, focusing primarily on the extensive margin, i.e., occupations disrupted and new opportunities created by AI. In contrast, this

study focuses on the intensive margin of workdays within the framework of a principal-agent model. Needless to say, we also build on and contribute to the literature that utilizes time allocation surveys, which have predominantly examined general or cyclical trends and their heterogeneity across population subgroups.[4] Among studies built on time allocation surveys, our study is unique in its focus on AI exposure, challenging the conventional expectation that technology frees humans from prolonged workdays.

The remainder of the paper is organized as follows: Section 2 develops a simple model within a principal-agent framework to provide theoretical guidance on the various ways AI technology can influence worker time allocation. Section 3 introduces the primary datasets used in our analyses, including patent data, occupation data, LinkedIn, Glassdoor, and the American Time Use Survey. Section 4 presents the empirical analysis on the relationship between AI exposure and work and leisure hours. Section 5 tests the mechanisms based on the principal-agent model. Finally, Section 6 concludes.

2. Modeling Framework and Hypotheses

Theories addressing the principal-agent problem have inspired a large body of research, including many seminal papers. While this study is primarily empirical, we ensure that our analyses are well-informed by theoretical insights. In particular, we build on straightforward adaptations of the Holmstrom and Milgrom (1987) model of dynamic incentive contracts, which examines how risk-averse agents respond to compensation schemes that balance incentives, risk-sharing, and the timing of information disclosure in a continuous-time framework. This model offers predictions about the relationship between a worker's "effort" (mapped into the number of work hours in our empirical context) and several key factors, including marginal productivity, the accuracy and timeliness of effort monitoring, and the worker's bargaining

[4]For instance, Aguiar et al. (2021) show that younger men experienced the greatest decline in market work hours among all demographic groups over the last 15 years, reallocating their leisure to video gaming and other recreational computer activities. Aguiar and Hurst (2007) find that the least educated adults experienced the largest increases in leisure. Aguiar et al. (2013) investigate how individuals reallocate their lost work hours during recessions. One exception is Ben-Rephael et al. (2025), which uses managers' time allocated to Bloomberg usage as a measure of effort provision and examines its' impact on firm value.

power in capturing or preserving the rents from technology-driven productivity gains.

The simple model, presented in the Online Appendix A, features a risk-neutral principal, a risk-averse agent, and a production output process following the standard Brownian motion in which effort and marginal productivity are multiplicative in determining the drift while noise is exogenously given. Under constant absolute risk aversion (CARA) utilities and a convex cost of effort for the agent, Holmstrom and Milgrom (1987) demonstrate that the optimal dynamic contract converges to a linear form in the aggregate: a lump-sum payment plus a share of the output, i.e., $\alpha + \beta X$. In this framework, the lump sum ensures the agent's reservation utility, $\underline{U}$ (shaped by the worker's relative bargaining power, which depends on the competitiveness of both the labor and product markets). The "sharecropping" coefficient, β, is inversely related to the agent's increasing marginal cost of effort, risk aversion, and output noise. Finally, the agent's effort level, in response to the incentive, is positively correlated with their marginal productivity and aligns in direction with the factors influencing β.

CARA utility abstracts from the wealth effect on leisure, a feature that may be unrealistic in many settings. The model can be extended to incorporate a general constant elasticity of substitution (CES) utility function, where the marginal utilities of consumption and leisure are interdependent. This framework allows for the examination of how the work-leisure allocation changes in response to external factors that influence the agent's reservation utility $\underline{U}$ via their best alternatives in the marketplace. When consumption and leisure are complements, that is, people enjoy consumption more when they have more leisure (MaCurdy, 1981; Blundell and MaCurdy, 1999); or when the reservation utility is sufficiently high (limiting the principal's ability to increase β due to the agent's risk aversion), work time is expected to decrease as the reservation utility rises. Since leisure is a normal good, the agent places greater value on it as their welfare improves. Rising $\underline{U}$ allows the agent to allocate more time to leisure and less to work (while enjoying higher consumption), all else being equal.

The model offers tight guidance on how AI can influence optimal incentives and the equilibrium level of effort for several reasons. First, if AI enhances the marginal productivity of

the agent, that is, if human and AI are complements in job tasks,[5] this increase in marginal productivity results in greater effort or longer working hours. Conversely, if human and AI are substitutes in the job,[6] the effect is reversed. It is worth noting that a principal-agent relationship is not required for this effect, as the same dynamic would apply to self-employed individuals.

Second, AI enhances work monitoring by providing better predictions or more precise signals of workers' efforts. This can occur through improved forecasting of market opportunities, ensuring that the right products are produced, or through more accurate assessment of workers' labor input using past and concurrent, own and peer data. Both mechanisms reduce the noise component (i.e., factors unrelated to workers' effort or actions), thereby increasing work hours. This effect operates in the same direction regardless of whether AI substitutes or complements labor, though it is significantly stronger when the worker acts as an agent (i.e., employed by someone else) rather than as a principal (i.e., self-employed).

Third, market forces and competitive conditions determine the extent to which workers benefit from AI-enabled productivity gains. When AI complements human labor and enhances labor productivity, the degree to which these gains translate into worker welfare–through a combination of higher pay and lower work hours—depends on the relative bargaining power of workers vis-a-vis their employers. Workers in regions or occupations characterized by competitive labor markets have limited bargaining power and may see little material benefit, with most of the rents accruing to employers or shareholders. Moreover, the share of rents available for firms to split with their workers also depends on product market competition. In highly competitive markets, consumers emerge as the primary beneficiaries of AI-driven productivity gains through better-quality products, lower prices, and rising consumer expectations, leaving little surplus for firms to share with their workers. If AI substitutes human labor and reduces

[5] A burgeoning literature corroborates complementarity in a wide range of occupations: lawyers(Armour et al., 2022), floor traders(Brogaard et al., 2024), stock analysts(Cao et al., 2024), and medical professionals(Wang et al., 2024).

[6] An equally large literature has expressed concerns over displacements of human labor and suppressing wages during technology advancement especially those targeted at routine-bases tasks: Kogan et al. (2023); Cheng et al. (2024); Hui et al. (2024);Jiang et al. (2025); Tuzel and Zhang (2021) and Zhang (2019).

labor productivity, workers find themselves in an even weaker bargaining position.

The distribution of the rents impacts work hours via the income effect linked to workers' reservation utility. When workers are able to capture a significant portion of the gains, their reservation utility increases, leading to greater consumption of leisure (a normal good), which, in turn, suppresses work hours. Conversely, when workers receive only a small share of the gains, the income effect from reservation utility is limited, resulting in minimal impact on work hours. The distribution of productivity rents serves as a distinct channel through which AI influences work-life balance.

3. Data, Measurement, and Overview

3.1. American Time Use Survey (ATUS)

The American Time Use Survey (ATUS), conducted by the Bureau of Labor Statistics from 2004 to 2023, is a primary resource for studying how people allocate their time. It provides comprehensive, nationally representative data on how Americans spend their time, where they spend it, and with whom. As the only federal dataset that captures both market activities (e.g., employment) and non-market activities (e.g., childcare, volunteering), ATUS has been widely used to investigate trends in work, leisure, health, and inequality (e.g., Aguiar et al., 2013, 2021; Alon et al., 2020; Doepke et al., 2023; Graff Zivin and Neidell, 2014; Krueger and Mueller, 2010).

The ATUS conducts a cross-sectional survey each year, with an average annual sample size of approximately 26,400 respondents. Participants are selected from households that have completed the Current Population Survey (CPS). Within each CPS household, one individual aged 15 or older is randomly chosen to complete the ATUS questionnaire. Following Aguiar et al. (2013), our sample consists of respondents aged between 16 and 65, excluding individuals who are not in a position to be employed, such as full-time students aged below 25 and those serving in the military.[7] As ATUS does not filter by employment status, unem-

[7]The military sector is defined using the Census industry code ("teio1icd") provided by ATUS, including

ployed respondents remain in our sample whenever an occupation code is available, generally reflecting their most recent jobs.[8] These criteria result in 131,324 unique individuals in the ATUS sample from 2004 to 2023. For the purpose of our research, our main analyses further exclude respondents from the technology sector so that we focus on workers in AI-using sectors rather than AI-inventing sectors.[9] With this exclusion we are left with 124,385 unique respondents.

A single interview is administered to each ATUS respondent by the Bureau of Labor Statistics, during which the prior day's activities are logged in a 24-hour diary segmented into 15-minute intervals. These activities, classified into over 400 distinct types, are grouped into four broad categories: basic survival (a fixed seven hours per day for critical survival functions such as sleeping and eating), market work (to be explained shortly), leisure, and others. Following previous literature (e.g., Aguiar et al., 2013, 2021; Boerma and Karabarbounis, 2021), our paper uses weekly hours as the unit of analysis, calculated by multiplying daily hours by seven (capped at 168 hours).

Market work, or simply "work," comprises main jobs, overtime work, work activities performed at home,[10] and supplementary tasks, such as security procedures and waiting related to work. "Work" time in our analysis encompasses the following ATUS-classified activities: "work, main job," "eating and drinking as part of job", "sports and exercise as part of job," "security procedures as part of job," "waiting associated with work-related activities," and "work-related activities, not elsewhere classified." Commuting and social activities at work are excluded (though including them yields qualitatively similar results).[11] In our empirical

national security and international affairs (9590) and armed forces (9600-9900).

[8]Long-term unemployed individuals do not have relevant occupation affiliation. Their exclusion does not impact our analysis of work time across occupations with varying AI exposure because they have no meaningful affiliation to any occupation.

[9]Following the literature (Acemoglu et al., 2022; Babina et al., 2024), the tech sector is defined using the Census industry code ("teio1icd") provided by ATUS, including information (6470–6780), scientific and technical Services (7380, 7460), and other professional, scientific, and technical services (7490). Details on this classification system can be found in Appendix A of the ATUS Data dictionary at https://www.bls.gov/tus/dictionaries/atusintcodebk23.pdf.

[10]Secondary jobs, if any, are excluded due to the lack of occupation-related information.

[11]Social activities at work include "socializing, relaxing, and leisure as part of job," and "travel related to work."

study, workday length is a proxy for effort provision—a common practice in the literature, e.g., Bandiera et al. (2020) for CEOs, Ben-Rephael et al. (2025) for executives, and Fehr and Goette (2007) for workers.

Leisure activities encompass activities such as watching television and movies, recreational computing and video games, reading, sports, and various hobbies. Since eating, sleeping, and personal care (ESP) fulfill essential biological needs and can also provide leisure value, any time beyond seven hours per day in these is thus counted as leisure. The residual category, "other," covers all remaining time, including home production (domestic responsibilities such as cleaning, maintenance, cooking, shopping, and gardening), childcare, education (personal academic pursuits, such as participating in classes or doing homework), job search activities (submitting resumes, conducting job interviews, and exploring employment opportunities), own medical care, civic activities (going to church or social club, volunteering, etc.), and any unclassified activities.

Panel A of Table 1 provides summary statistics at the ATUS respondent level. Unless otherwise specified, all potentially unbounded variables are winsorized at the 1% extremes. The average respondent allocates 35.2 hours to work and 55.3 hours to leisure per week. The variation is substantial, with standard deviations of 30.4 and 27.2 hours, respectively. Within the residual category, the average respondent spends 1.2 hours on education, 1.6 hours on civic activities, 0.4 hours on own medical care, 0.1 hours on job search, 15.9 hours on home production, and 4.6 hours on child care. These time allocation estimates are consistent with earlier studies (e.g., Aguiar et al., 2013, 2021).

ATUS also reports wages for each individual, which are converted into 2023 constant dollars in our analysis. For hourly workers, the hourly wage is directly reported; for non-hourly workers, we estimate the hourly wage as their weekly earnings divided by what the respondents self-report as their "usual" work hours per week. About 37.5% of the respondents report a usual workweek of 40 hours, 30.8% report more than 40 hours, and 31.6% report fewer than 40 hours. The average hourly earnings in our sample are $28.1 (in 2023 dollars).

[Insert Table 1 here.]

3.2. AI patents

Central to our analysis is quantifying individual occupation's exposure to AI technologies. Following recent literature on technology disruptions, we use the textual correlation between AI patents and job task descriptions as our AI exposure measure. The first step is thus to collect a comprehensive sample of AI patents granted between 2000 and 2023 from the Artificial Intelligence Patent Dataset (AIPD). AIPD was first publicly released by the United States Patent and Trademark Office (USPTO) in 2021 and expanded in 2024 to include all patent documents published through 2023. Pairolero et al. (2025) detail a machine learning procedure, adopted by AIPD, that assigns each U.S. patent (1976–2023) a probability of being AI-related. Patents are classified as AI patents if that probability exceeds one of the three thresholds: 50%, 86%, or 93%. Pairolero et al. (2025) suggest the 86% threshold as the best trade-off between precision (correctly identifying AI patents) and recall (capturing the full set of AI patents); we therefore adopt this cutoff.

These procedures yield a total of 905,667 AI patents granted between 2000–2023, which are classified into one or more of the eight categories defined by Pairolero et al. (2025): (i) machine learning, (ii) vision, (iii) natural language processing, (iv) speech, (v) evolutionary computation, (vi) AI hardware, (vii) knowledge processing, and (viii) planning and control.

Prior research indicates that only a small subset of patents have meaningful scientific and economic value. For example, about one quarter of patents remain uncited, and fewer than one percent receive more than one hundred citations (Kogan et al., 2017). To focus on technologies with the greatest transformative potential, we limit our analysis to the top 1% of AI patents each year, identified by their forward citation counts adjusted for both technology class and vintage. Following Kogan et al. (2017), an adjusted forward citation count is calculated by dividing each patent's raw citation count by the average citation count of AI patents granted in the same CPC subclass and year-quarter. This selection yields a final sample of 9,270 AI

patents. Online Appendix Figure OA.1 shows the annual number of these patents and their mean adjusted citation counts.

The textual information from the title and abstract of each AI patent allows us to extract information about the scope and content of the underlying innovations from text corpora. This information is then matched to occupations to assess the latter's exposure to AI.

Figure 1 presents three word clouds that trace the evolving landscape of AI patents in our sample over time. From 2000 to 2009, AI innovation centered on smoothing human input and managing data, with dominant keywords including "computer," "information," "interface," and "image." Between 2010 and 2019, emphasis shifted toward "security," "monitoring," "voice," and the early emergence of "generate." Since 2020, the word cloud has expanded to include "virtual," "generate," "video," "autonomous vehicle," and "automate," reflecting AI's transformation from data and interface tools to immersive perception, autonomous operation, and creative synthesis.

[Insert Figure 1 here.]

Figure 2 presents the timeline of ten high-impact AI patents in our sample, selected to be roughly evenly distributed across the 2000–2023 period. Each patent on this list was chosen based on a combination of adjusted citation counts, technological influence, and representation of key AI application areas. Early developments are exemplified by Qualcomm's handwriting annotation (2000) and Microsoft's auto-completion (2002), which reflect AI's initial focus on enhancing user input and productivity tools. The mid-2010s highlight a transition toward richer interaction and sensor-driven services, as illustrated by Meta's personalized feed (2010), Apple's multi-touch gestures (2013), and Skybell's doorbell communication systems (2015). More recent patents reflect the frontier of AI capabilities, such as Nvidia's real-time lane detection (2021) and Google's generative search summaries (2023), signaling AI's move into real-time perception and content synthesis.

[Insert Figure 2 here.]

3.3. Occupation data

The second step in completing the measurement involves retrieving job tasks from U.S. Department of Labor's Occupational Information Network (O*NET), which classifies each occupation by an 8-digit Standard Occupational Classification (SOC) code and provides detailed descriptions of its specific tasks. For example, in 2023 the Data Scientist occupation (SOC 15-2051.00) includes tasks such as "analyze, manipulate, or process large sets of data using statistical software," "create graphs, charts, or other visualizations to convey the results of data analysis using specialized software," and "propose solutions in engineering, the sciences, and other fields using mathematical theories and techniques."

Every year, O*NET covers 900–1,100 occupations, each identified by an 8-digit SOC code with detailed job task information. We draw on every historical release of the O*NET database to create an annual panel of occupations from 2000 to 2023. Section B.1 of Online Appendix provides more details of the procedure.

3.4. Supplemental data on employment: LinkedIn and Glassdoor

Two databases provide supplementary information on employment, workload, compensation, and job satisfaction at the individual level. LinkedIn dataset from Revelio Labs supplies structured, resume-style profiles detailing salary and employment history. O*NET covers approximately 800 6-digit SOC occupations, of which the LinkedIn data includes 335—each mapped to the Census occupation codes used in ATUS (see section 3.6.3). This structure supports aggregation at both the occupation and employer levels. Using employment histories through mid-2023, we aggregate individual data at the occupation × firm × year level and assemble firm × year panels to investigate the relationship between time allocation and firm outcomes.[12]

Glassdoor data, also accessed via Revelio Labs, provide detailed information about pay, workload, and employee reviews for a broad set of firms, including all major employers. Prior

[12]We exclude observations missing occupation or firm information and drop firms with fewer than 100 US-based employees in the prior year to avoid noise from poor coverage following Fedyk and Hodson (2023).

work shows that the Glassdoor data provide valuable insights into firm performance and labor market dynamics, though they tend to overrepresent skilled occupations (e.g., Edmans, 2011; Green et al., 2019; Gornall et al., 2024). Each employee review contains text and ratings (ranging from 1-5) on multiple dimensions, including two metrics that are most relevant to our study—overall satisfaction and work-life-balance (WLB) at the firm. Reviewer information includes job title, tenure, employment status, and location. Limiting the sample to US-based employees and firms with a minimum of 20 reviews in a given year yields 2,607,571 reviews across 3,869 firms.

Panel B of Table 1 reports the summary statistics of the employee rating sample at the occupation × firm × year level. Each cohort has an average of 4.75 reviews. The average overall job rating is 3.50 (out of 5.0), and 3.40 for work-life balance. The annual salary averages $87,710, with a standard deviation of $44,240.

3.5. Other data

For firm performance analyses, the sample is limited to U.S. publicly listed companies using Compustat, CRSP, and other related WRDS databases. [13] *ROA* (return on assets) is defined as the ratio of operating income before depreciation to total assets, with an average value of 5.4% and an inter-quartile range of 3.0% to 14.8%. Labor productivity is defined as sales over employment (in $000), and the average value is 1.95.

3.6. Measuring occupational AI exposure over time

3.6.1. Measuring AI exposure at the occupation-patent level

Accurately measuring occupational exposure to AI technology is essential for attributing changes in time allocation to AI. There have been a variety of exposure measures to a diverse set of technologies or innovations, e.g., AI exposure developed by Felten et al. (2018) and Webb (2019), generative AI exposure from Eisfeldt et al. (2023) and Hartley et al. (2024),

[13] Utility & finance sectors are excluded.

software and robot exposure developed by Webb (2019), fintech exposure from Jiang et al. (2025), labor-saving and labor-augmenting technology exposure from Kogan et al. (2023). These measures typically analyze the micro-foundations of tasks and aggregate each task's exposure to the occupational level based on task importance. Ideally, the exposure measure captures both cross-sectional differences across occupations and time-series variations in AI exposure within each occupation.

The measurement of occupational AI exposure in this study builds directly on two established methodologies. The first relies on the textual similarity between AI patents and job task descriptions; its principal advantage lies in the interpretability and determinacy in the resulting exposure measure. The second method leverages large language models (LLMs), such as ChatGPT, to extract and interpret relevant information from unstructured text. By virtue of their generative capabilities, LLMs offer greater flexibility in how tasks can be framed, better contextual understanding, and more nuanced language interpretation.

This study uses OpenAI's GPT model for AI exposure classification, as this represents the most recent and potentially powerful tool for text interpretation. Specifically, for each occupation (o) in year (t), we submit the complete set of that occupation's task descriptions alongside the text of every AI patent (i) granted in the same year to the GPT model, requesting a similarity comparison. To ensure robustness, we replicate our main results using a conventional textual similarity measure based on Term Frequency–Inverse Document Frequency (TF-IDF) at the same occupation-patent granularity. TF-IDF has been validated in prior studies as an effective and robust natural language processing (NLP) approach that weights terms by their specificity to individual documents relative to the broader corpus. In both approaches, matching patents and occupational task descriptions within the same year ensures that our measures track the time-varying content of each occupation.

Our sample comprises 9,064 high-impact AI patents granted between 2000 and 2023; the number of occupations averages about 950 each year, identified with the 8-digit SOC code. As a result, the GPT model encodes 8.6 million pairs at the occupation-year $(o, t) \times$ patent (i) level, yielding two key intermediate variables. The first, AI exposure score ($AI^{EXP}_{o,i,t}$), is a

correlation score (ranging from 1 to 10) between the text description of the title and abstract of patent i granted in year t and the full set of task descriptions for occupation o in the same year. Across all pairings, the mean correlation score is 3.7 and the standard deviation is 1.8.

The second intermediate variable is a complementarity classification ($AI_{o,i,t}^{COMP}$) following procedures of Kogan et al. (2023) and Jiang et al. (2025). It is a categorical variable (1 = complement, 0 = neutral, and −1 = substitute) that indicates whether a given AI patent primarily complements, substitutes, or is neutral to the tasks of an occupation. Among all occupation-patent pairs, 77.4% are classified as complementary, 19.4% as substitutive, and 3.2% as natural.[14]

Section B.3.1 of the Online Appendix provides additional details on the GPT prompt setup, examples, and validation. Importantly, our analysis relies not on the absolute scale of the two variables, but solely on their relative values to position the AI exposure of occupation-patent pairs in both the time series and the cross section.

3.6.2. Aggregating AI exposure to the occupation-year level

To measure the aggregate impact of a cluster of AI innovations on a given occupation, we sum up the impact of individual AI patents published during the previous five-year period leading to the current year. That is,

$$AI_{o,t}^{EXP} = \sum_{i \in \Theta_t} AI_{o,i,t}^{EXP}, \tag{1}$$

where Θ_t represents the set of all AI patents i published between year $t-4$ and year t.

It is natural to compare the resulting measure with other measures of AI exposure in the existing literature, particularly those from Webb (2019), Felten et al. (2019) and Hampole et al. (2025). Webb (2019) applies natural language processing algorithms to measure the overlap between text descriptions of job tasks and patents. Felten et al. (2019) link the workplace abilities of occupations to the progress of nine AI applications (such as speech recognition and

[14]This proportion is almost identical to Kogan et al. (2023)'s finding that approximately 19.7% of the job tasks are susceptible to AI substitution.

image generation) tracked by the Frontier Foundation (EFF) from 2010 to 2015 using survey responses. There are two main differences between our measure and these two earlier ones: First, both previous measures are time-invariant and are based on information at the end of their respective sample periods. Second, due to their research focus and the sample periods ending in mid- to late- 2010s, AI technologies have evolved significantly in preexisting classes and brought about two new (out of the eight) AI technology classes present in our sample.

Hampole et al. (2025) extract AI-related mentions from LinkedIn profiles and job postings and map those texts onto occupational task descriptions. By contrast, our measures build on a growing literature that compares the textual content of AI-related patents to occupational task descriptions. Moreover, the later approach allows us to fully explore the approximately 800 occupations covered by O*NET (at the SOC 6-digit level) for two decades, as opposed to the 335 SOC 6-digit codes covered by Revelio Labs.

Textual similarity-based measures of technology exposure are inherently non-directional; they do not distinguish whether AI substitutes or complements labor. Empirical studies have presented mixed findings regarding the directional effect of AI exposure. Some report labor-displacing effects of AI: for example, (Hampole et al., 2025) find that occupations highly exposed to AI tend to experience reduced labor demand, although productivity gains boost overall employment across the broader labor market. Others find complementary effects, such as Liu et al. (2023), who document that AI exposure is associated with increased job postings. Still others highlight heterogeneous effects: Berger et al. (2024) show that Generative AI tends to complement high-level white-collar jobs while substituting for lower-level ones.

For this reason, our empirical tests build on an ex ante decomposition of the substitutive and complementary effects of AI exposure. The construction of an AI net complementarity exposure follows the approach used in Jiang et al. (2025) for the context of fintech. Specifically, for a given SOC 8-digit occupation o in a year t, AI net complementarity ($AI^{COMP}_{o,t}$) is defined as the sum of the product of AI exposure and AI complementarity classification (with value from $\{-1, 0, +1\}$) of occupation o with respect to AI patents i published during the five-year period ending in year t, as shown in the following equation:

$$AI_{o,t}^{COMP} = \sum_{i \in \Theta_t} AI_{o,i,t}^{EXP} \cdot AI_{o,i,t}^{COMP}. \tag{2}$$

For ease of interpretation, we normalize both $AI_{o,t}^{EXP}$ and $AI_{o,t}^{COMP}$ by dividing each by 10,000. This scaling ensures that the typical exposure value falls between 0 and 1. The resulting measures have mean values of 0.56 and 0.39, and standard deviations of 0.41 and 0.38, respectively.

3.6.3. Matching occupation-level AI exposure to the ATUS respondents

The ATUS data use Census occupation classification codes, which must be bridged to our 8-digit SOC-based AI exposure measures. We therefore adopt the "*occ*1990*dd*," classification system developed by Dorn (2009) and its various updates to aggregate Census occupation codes into a balanced panel of occupations. The $AI_{o,t}^{EXP}$ and $AI_{o,t}^{COMP}$ measures are then merged between *occ*1990*dd* occupation and SOC 6-digit occupation.[15]

Table 1 Panel A reports the summary statistics of the occupational AI exposure measures of ATUS respondents. The average $AI_{o,t}^{EXP}$ and $AI_{o,t}^{COMP}$ scores are 0.66 and 0.48, respectively. The positivity of $AI_{o,t}^{COMP}$ and its proximity to $AI_{o,t}^{EXP}$ in magnitude indicates that AI innovations tend to have a complementary effect rather than a substitutive effect on the labor market. Panel A of Figure 3 shows the time series of the average AI exposures, $AI_{o,t}^{EXP}$ and $AI_{o,t}^{COMP}$, whichhich summarize the occupations held by ATUS respondents from 2004 to 2023. Both exposures experience a four-fold increase during the sample period.

[Insert Table 1 here.]

[15]The "*occ*1990*dd*" classification system has been widely employed in labor economics studies (e.g., Autor and Dorn, 2009, 2013; Webb, 2019). Crosswalk and documentation are available at https://www.ddorn.net/data.htm. We match the SOC 6-digit occupation codes to *occ*1990*dd* in three steps: (i) we first match SOC 2000 codes and SOC 2018 codes to SOC 2010 codes using crosswalks provided by BLS at https://www.bls.gov/soc/soc_2000_to_2010_crosswalk.xls and https://www.bls.gov/soc/2018/soc_2010_to_2018_crosswalk.xlsx; (ii) We then use crosswalk provided by Webb (2019) to map the SOC 2010 codes to the 2010 Census occupation codes; (iii) lastly, the 2010 Census occupation codes is matched to *occ*1990*dd* codes using crosswalk provided by Autor (2015) at https://www.ddorn.net/data.htm. These three steps address modifications in SOC 2018 (the last major update since 2018) by matching new occupation codes with *occ*1990*dd* if (1) the SOC 2018 codes can be mapped to SOC 2010, and (2) the corresponding SOC 2010 codes can be linked to 2010 Census occupation codes. Sporadic and minor changes since 2018 are classified manually.

[Insert Figure 3 here.]

Finally, the raw scores of the AI exposure measures are transformed into percentile ranks (where 1 and 100 represent the lower and upper bounds) each year, following the literature (e.g., Autor and Dorn, 2013; Webb, 2019). For the interest of the readers, Table OA.2 of the Online Appendix lists the top occupations sorted by AI exposure and AI net complementarity in 2023. At the top of the list of both $AI_{o,t}^{EXP}$ and $AI_{o,t}^{COMP}$ are computer and information system managers, bioinformatics technicians, operations research analysts, and management analysts. Occupations with high $AI_{o,t}^{EXP}$ but low $AI_{o,t}^{COMP}$ include data entry keyers, tellers, and office machine operators. Those at the bottom are common in both dimensions, including occupations of dancers, barbers, and meat packers.

3.6.4. Validation of occupation-level AI exposure measures

Generative AI has rapidly emerged as a powerful tool for information classification. Given its relatively short track record and evolving capacities, it is important to validate our GPT-generated AI exposure measure ($AI_{o,t}^{EXP}$) by comparing it with one derived from conventional machine learning techniques. Specifically, we benchmark our measure against a TF-IDF cosine similarity measure because of its intuitive term-weighting, transparency, and ease of replication. Section B.2.3 of the Online Appendix describes the full TF-IDF procedures. At the occupation-year level, the two measures exhibit a correlation of 0.51, reflecting moderate alignment. Moreover, substituting the TF-IDF measure for $AI_{o,t}^{EXP}$ in our main regressions yields qualitatively similar results.

The TF-IDF measure is not directional and therefore cannot help validating our AI net complementarity exposure ($AI_{o,t}^{COMP}$). Instead, we perform an indirect validation by relating $AI_{o,t}^{COMP}$ to the wage growth predictions due to AI complementarity and substitution, estimated in Kogan et al. (2023).[16] In 2023, $AI_{o,t}^{COMP}$ exhibits a correlation of 0.60 with the overall

[16]Using an open question-based approach, Kogan et al. (2023) ask ChatGPT about AI's potential to substitute or complement job tasks and yields time-invariant measures of different AI exposure components. They do not report the exposure but provide AI-related earnings changes of occupations with the highest complementarity (substitution) exposure in Online Appendix. Section B.2.4 of the Online Appendix provides more

wage growth related to AI. When decomposed, the wage growth components attributed to AI substitution and complementarity correlate with $AI_{o,t}^{COMP}$ at -0.59 and 0.47, respectively. Taken together, this evidence validates the reliability of our AI net complementarity measure.

3.6.5. Comparison with other occupation-level exposure measures

A growing literature has estimated and analyzed occupation exposure to a wide array of technologies and innovations, including AI. It is thus obligatory for us to compare and distinguish AI exposure from the other exposure measures. Figure OA.2 of Online Appendix plots the time-compressed version of our AI exposure (by averaging over the years from 2000–2023), against six time-invariant occupational exposure measures developed in earlier studies: AI exposure by Felten et al. (2019), AI exposure and robot exposure by Webb (2019), routine task intensity (RTI) by Autor and Dorn (2013), offshorability exposure developed by Firpo et al. (2011) and standardized by Autor and Dorn (2013), and work-from-home (WFH) feasibility score by Dingel and Neiman (2020).

Figure OA.2 in the online appendix shows that our and the two other AI exposure measures (Felten et al., 2019; Webb, 2019) are positively correlated. Its relations with robot (Webb, 2019) and RTI exposure (Autor and Dorn, 2013) are not monotonic. Finally, our AI exposure is positively correlated with offshorability and WFH potentials, but only at the high percentile levels.

4. AI and Workday: Empirical Relations

4.1. Event study: ChatGPT

The release of generative AI tools, notably ChatGPT in November 2022, marked a watershed moment for AI adoption in the workplace. Its immediate accessibility and versatility accelerated AI integration across industries and transformed business processes almost overnight. According to McKinsey (2024), 33% of respondents' organizations adopt generative AI right

details on the validation procedures.

away in 2023 and that number increases to 65% in 2024.[17] The advent of generative AI was a transformative event—rather than a gradual progression—whose precise timing was unforeseen by its adopters. Such properties make it a desired setting to study the impact of AI on workday.

The event study entails a difference-in-differences framework around the shock, with time allocation variables from the ATUS data as the dependent variables. The test sample covers 2022 to 2023, a relatively short range striding ChatGPT to capture the discrete change without being confounded by longer-term trends. The premise of the test is that the impact of AI adoption on work hours should be more prominent among occupations with greater sensitivity to generative AI. In other words, the level of "treatment" is captured by the generative AI exposure of the occupation to which a worker is affiliated, as defined by Eisfeldt et al. (2023).[18] The regression, at the survey respondent level, with subscripts of i (individual), o (occupation), and t (year) level, is as follows:

$$Y_{i,o,t} = \beta_1 \cdot GenAI_o^{EXP} \cdot POST_t + \beta_2 \cdot X_{i,t} + \alpha + \epsilon_{i,o,t}. \tag{3}$$

The dependent variable is the number of weekly hours allocated to each activity category (i.e., market work or leisure). $GenAI_o^{EXP}$ is generative AI exposure for occupation o, measured by its percentile rank. The $POST_t$ dummy equals one for the year 2023. The regression incorporates a set of individual-level controls, including age, the number of children below 18, and a set of indicators for gender, educational attainment, marital status, and race. The regression further includes a battery of fixed effects, α, at the following levels: occupation, state $\times$ year, industry $\times$ year, year-month, and day-of-week.[19] These fixed effects filter out macroeconomic factors at both the industry and state levels, as well as seasonality and weekday

[17] https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai

[18] Eisfeldt et al. (2023) use a large language model to classify whether job tasks of occupations can be performed more effectively using ChatGPT based on task descriptions. Data is available at https://sites.google.com/view/gregorschubert/data?authuser=0.

[19] Industry is defined by the Census detailed industry code "trdtind1" used in ATUS (ATUS identifies 51 unique industries; see Appendix A of the ATUS Data dictionary at https://www.bls.gov/tus/dictionaries/atusintcodebk23.pdf).

effects.

In all regressions throughout this study, we adhere to the following best practices. First, all linear regressions using ATUS data are weighted by ATUS sample weights in order to recover the representativeness of the population (e.g., Aguiar et al., 2021). Second, standard errors are double-clustered at the occupation and state levels. Third, unless otherwise noted, all potentially unbounded variables are winsorized at the 1% extremes.

Table 2 reports the weighted linear regression results for equation (3) using ATUS data from 2022 to 2023. Column (1) reports the results for the full sample. Specifically, workers more exposed to generative AI experienced significantly increased work hours (Panel A) and reduced leisure hours (Panel B) following the introduction of ChatGPT. Comparing 2023 to 2022, an interquartile increase in generative AI exposure corresponds to an additional 3.75 hours (0.075×50 percentiles from the 25th to the 75th) of work and a reduction of 3.85 hours in leisure. Such a magnitude is economically significant as they represent 10.65% and 6.97% of the average work and leisure hours, respectively.

[Insert Table 2 here.]

Columns (2) and (3) present the results for subsamples divided into the top quartile and the remaining observations based on the extent to which generative AI complements the job tasks (see definition in Section B.3 of the Online Appendix). For work hours, the coefficient magnitudes for the top-quartile subsample are approximately twice those of the remaining observations, and about 50% larger for leisure hours, although these differences are not statistically significant. Columns (4) and (5) present results for subsamples divided into the top quartile and the remaining observations based on local AI awareness, measured by state-level Google search trends for ChatGPT from November 30 to December 31, 2022.[20] Workers in regions within the top quartile of AI awareness experience greater association between generative AI exposure and work and leisure hours than the remaining observations.

[20]Figure OA.3 of the Online Appendix plots the Google search trend of AI and ChatGPT from 2010 to 2023, confirming a peak in Google search of not just "ChatGPT" but also "AI" in December 2022 following the release of ChatGPT.

In addition to the average impact, a potentially interesting question is whether Gen-AI makes the "already over-worked" individuals work more or less. For the empirical test, we define an "already overworking" subsample as workers from occupations whose average reported normal work hours per week is no fewer than 45 hours in 2021 (about 11.7% of the respondents). We find that the positive (negative) relationship between generative AI exposure and work (leisure) hours appears only among workers who are not yet overworking. This suggests that AI primarily affects those with more manageable schedules, pushing them closer to their limits, rather than exacerbating workloads for those already burdened.

4.2. Occupation AI exposure and workday

4.2.1. AI exposure and workday in the long panel

Next, we extend the event study to the full sample period using measured occupational AI exposure (see Section 3.6 for details). As the first step, Figure 4 serves as a diagnostic test, presenting the scatter plots and quadratic-fitted lines for the association between occupational AI exposure and work hours in the first (2004-2013) and last half (2014-2023) of our sample period. Individuals in higher occupational AI exposure report longer work hours. Notably, the quadratic-fitted line for the 2014-2023 consistently lies above the line for the 2004-2013 with steeper slopes, indicating a stronger impact of AI on work-life balance in recent years.

[Insert Figure 4 here.]

For the formal analysis, we estimate the relationship between occupational AI exposure and work and leisure hours at the individual (i) respondent level, indexed by occupation (o) and year (t)s:

$$Y_{i,o,t} = \beta_1 \cdot AI^{EXP}_{o,t-1} + \beta_2 \cdot X_{i,t} + \alpha + \epsilon_{i,o,t}. \tag{4}$$

where the dependent variables are weekly hours spent on market work and leisure. The key

independent variable, $AI^{EXP}_{o,t-1}$, is the lagged occupational AI exposure, constructed as detailed in Section 3.6. Otherwise, the model specification, including control variables for individual characteristics, fixed effects, and standard errors, follows the same as in Equation (3).

Table 3 reports the weighted linear regression results for Equation (4) using the ATUS data from 2004 to 2023. Columns (1)–(3) present the results for weekly work hours. Higher occupational exposure to AI is associated with increased work hours. Specifically, in column (1), where all fixed effects are included except the occupation fixed effects, an interquartile increase in occupational AI exposure increases workday by 2.25 ($= 0.045 \times 50$) hours per week on average. The cross-sectional relationship is both economically and statistically significant (at the 1% level).

To alleviate concerns regarding confounding factors, the regression in column (2) further controls for other common occupation exposure measures in the literature, including robot exposure by Webb (2019), routine task index (RTI) by Autor and Dorn (2013), and offshorability exposure constructed by Firpo et al. (2011) and standardized by Autor and Dorn (2013), all in percentile ranks. The direction and magnitude of the coefficient for occupational AI exposure remains consistent and significant at the 1% level. To further mitigate the concern that the relation between AI exposure and time allocation could be driven by occupation-level unobserved heterogeneity, column (3) incorporates occupation fixed effects. The magnitude of the coefficient is about one-half that of column (1), implying an interquatile effect of 1.15 hours, which remains significant at the 5% level.

[Insert Table 3 here.]

Columns (4)–(6) of Table 3 report the results for weekly leisure hours. Leisure hours decline as occupational AI exposure increases. Specifically, column (4) shows that an interquartile increase in AI exposure is associated with a 1.55-hour reduction ($= -0.031 \times 50$) in weekly leisure time on average. This negative effect remains robust when additional occupational exposure measures are included (column (5)) and when occupation fixed effects are added (column (6)). The combined results also suggest a slight decrease in time allocated to the

residual category (which includes personal care, education, and related activities), helping to accommodate the remaining gap between work and leisure hours. Given the mostly symmetric effect between work and leisure, we focus on work hours alone for the remaining analyses.

Table OA.3 of the Online Appendix reports the results of a battery of robustness tests using alternative model specifications. Column (1) uses an alternative AI-exposure measure, defined as the percentile rank of similarity scores between TF–IDF representations of job-task descriptions and AI patents granted over the five-year period ending in the current year.[21] In column (2), the dependent variable of work hours is modified to include time spent on commute, work-related travel, and social and leisure activities at work. The rest of the table addresses multiple specificities, including part-time status (column (3)); control for lagged usual work hours (column (4)); exclude unemployed (columns (5)–(6)) and absent (currently employed but absent from work on the survey date) individuals (columns (7)) or weekends (columns (8)), and isolate the subsample of workers who are compensated on an hourly basis (and thus command greater flexibility in adjusting work hours) (columns (9)). The estimated effects of occupational AI exposure on work hours remain significantly positive and align closely with the results in Table 3.

4.2.2. Decomposition of leisure activities

Given that leisure is a large collection of diverse activities, a natural question arises: what kinds of leisure activities give way to longer work hours? Table 4 explores this question by decomposing total leisure into screen-based and non-screen-based activities. This distinction is not merely descriptive; it reflects fundamental differences in the cognitive and restorative nature of leisure. Screen-based activities, such as recreational computer use, video gaming, and watching TV, are often more passive and less physically or socially engaging. In contrast, non–screen-based activities—reading, sports, listening to music, and travel—are typically more active, effortful, or immersive, and are associated with greater psychological restoration and well-being (de la Rosa et al., 2024).

[21] Section B.2.3 of the Online Appendix provides detailed descriptions for the scores based on TF–IDF.

[Insert Table 4 here.]

Columns (1) and (2) of Table 4 indicate that the previously observed decline in total leisure time is primarily driven by a reduction in non-screen-based activities, while time spent on screen-based leisure remains largely unaffected by occupational AI exposure. Columns (3) through (6) further break down non-screen-based activities into four categories: recreation (e.g., relaxing, listening to music, traveling), socializing, the leisure components of eating, sleeping, and personal care (ESP), and other activities (e.g., hobbies, reading, sports). This shift is concerning because time is being drawn from activities crucial for recovery, especially given that screen-based leisure—less effective at restoring cognitive and physical energy—remain unchanged.

4.2.3. Dispersion of workday in relation to AI

The positive relation between AI exposure and length of workday does not provide information about within-group dispersion. This section explores dispersions along three dimensions.[22]

The first dimension of dispersion we examine is gender. The heterogeneous impact of AI with respect to gender is unclear a priori. On the one hand, studies finds that automation can reduce gender gap in employment and wages (e.g., David and Melanie, 2013; Acemoglu and Restrepo, 2022; Cortés et al., 2024). On the other hand, Cook et al. (2021) show that women's higher opportunity cost of time in the form of unpaid work may sustain the gap in gig economy settings. The second dimension is age. Technologies, even when labor augmenting, can widen pay disparities by disproportionately benefiting younger workers who have the skills or flexibility to adapt (Kogan et al., 2023).

Table 5 presents regression at the occupation-year level where the dependent variable is the within-occupation workweek gap. Columns (1) and (2), which report "female minus male" differences, show that both general AI exposure and AI net complementarity are associated with a disproportionate increase in women's work hours relative to men's (significant at 5%

[22] An unsorted, simple analysis of a relation between within-occupation standard deviation of workday vs. AI exposure yields null results, see Table Table OA.4 in the Online Appendix.

level), thereby narrowing the gender gap, as women typically work fewer hours. Columns (3) and (4), which report "young minus old" differences, indicate that AI exposure increases work hours more for younger workers than for older ones (significant at 10% level), also reducing age-based disparities, given that younger individuals are typically less fully employed.

[Insert Table 5 here.]

The third dimension is within-household dispersion. Time allocation could be a decision made jointly by household members (Becker, 1965; Del Boca and Flinn, 2012). For instance, individuals with a busy spouse may need to dedicate more time to home production and child care. The ATUS does not link members of the same households; however, a subsample of ATUS respondents (about 59.4% of the full sample) report their spouses' employment status including the latter's typical workweek, providing us a way to explore the joint work-hour decision by interacting AI exposure with spouse employment status.

Table OA.5 in the Online Appendix reports the results. Column (1) replicates full sample baseline result (Table 3) on the subsample, confirming that general AI exposure is associated with longer work hours in this group. Column (2) includes an indicator, *I(Spouse Employed)*, which equals one if a respondent's spouse is employed in a given year and zero otherwise. A significantly negative coefficient in *I(Spouse Employed)* is consistent with the hypothesis of joint household allocation, but its inclusion does not alter the coefficient of AI^{EXP}. Columns (3) and (4) further segment the sample based on the spouse's usual weekly work hours. For respondents whose partners fall in the top quartile, the coefficient on $AI^{EXP}_{o,t-1}$ is negative, while it remains positive for all others, suggesting that the effect of AI on an individual's workday is tempered when their spouse is very time-constrained. Overall, the results provide some evidence of within-household reallocation of work hours in response to AI exposure.

4.2.4. Additional evidence: Residual categories and possibility of downsizing

Next, we investigate how occupational AI exposure influences time allocation in the residual categories, i.e., activities other than work or leisure. Table OA.6 in the Online Appendix sum-

marizes the findings. Most of the categories exhibit insignificant relations with AI exposure. The exception is civic activities which have a positive relation with AI exposure (significant at 5%). Hence, AI has not been a contributor to the secular decline in devotion to social work and community engagement, a phenomenon known as "bowling alone."[23]

Finally, we examine whether the lengthened workday is a by-product of AI's negative impact on employment, specifically, whether remaining workers are absorbing tasks from displaced colleagues. Table OA.7 reports a muted relationship between AI exposure and employment trends (both level and changes) from 2004–2023, based on the Occupational Employment and Wage Statistics (OEWS) data from the Bureau of Labor Statistics, consistent with prior research (e.g., Acemoglu et al., 2022; Hampole et al., 2025). While 2024–25 saw notable layoffs in the tech sector, our sample explicitly excludes this industry, and aggregate unemployment has remained relatively stable into 2025. Taken together, these findings indicate that extended workday is unlikely a result from firm downsizing or task consolidation among fewer employees.

5. Testing model predictions

5.1. Marginal productivity: AI complementarity vs. substitution

5.1.1. Technology complementarity

Technology can influence labor in two primary ways: substitution, where it replaces job tasks, and complementarity, where advancements in capital—such as improved tools—enhance workers' marginal productivity (e.g., Acemoglu, 1998; Acemoglu and David, 2011; Acemoglu and Restrepo, 2019). Thus, the overall effect of AI exposure shown in Table 3 invites a bifurcation into complementarity and substitution. To decompose general AI exposure, we use the GPT model to classify each AI patent as complementary, substitutive, or neutral to each task of an occupation based on its textual descriptions. AI net complementarity exposure at

[23]The term was coined by the book *Bowling Alone: The Collapse and Revival of American Community* (2000) by Robert D. Putnam.

the occupation level, $AI_{o,t}^{COMP}$, is defined as the difference between exposure to complementary and substitutive AI patents over the past five years, transformed into percentile ranks by year. That is, we rank occupations by their net complementarity scores to AI technology (with low complementarity indicating a strong substitution effect). Section 3.6 provides further details on variable construction.

Table 6 presents the weighted linear regression results for the impact of AI net complementarity exposure based on Equation (4), replacing $AI_{o,t-1}^{EXP}$ with $AI_{o,t-1}^{COMP}$. The dependent variables are weekly work hours in columns (1)–(3). Column (1) shows that, controlling for individual characteristics and fixed effects at the levels of state × year, industry × year, year-month and day-of-week, an interquartile increase in AI net complementarity is associated with an additional 2.8 ($= 0.056 \times 50$) work hours per week, equivalent to 7.95% of the sample mean (35.2 hours). Such a positive relationship between AI net complementarity exposure and work hours remains consistent when additional occupational exposure measures, including robot exposure, RTI, and offshorability, are included (column (2)); or with occupation fixed effects incorporated (column (3), significant at the 5% level).

[Insert Table 6 here.]

Overall, the magnitude and significance of the coefficients on $AI_{o,t-1}^{COMP}$ are greater than those of $AI_{o,t-1}^{EXP}$ presented in Table 3. The evidence suggests that the extended workday could be attributed to AI's complementarity to human work. In other words, people end up having longer workdays precisely when AI makes them more productive (and presumably saves them time on given tasks). The seeming paradox echoes 19th century economist Jevons (1865), who predicted that improvements in engine technology—and hence energy efficiency—would lead to increased demand for and consumption of energy (coal at the time). Labor is another factor of production that could apply the logic: When task productivity improves, the demand for additional tasks increases with AI, along with heightened expectations for both quality and expediency.

5.1.2. Wage effects

One might argue that workers are compelled to work longer hours to remain competitive, particularly when AI exposure induces a substitution effect. The distinction between substitution and complementarity can be examined by analyzing the relationship between AI exposure and wages. If AI complements labor by enhancing worker productivity, we should observe an increase in hourly wages (holding market competition constant); if it substitutes for labor, wages should stagnate if not decline. To test the two competing hypotheses, we re-estimate Equation (4) using wages from the ATUS, defined as 100 times the natural logarithm of hourly wages in 2023 constant dollars as the dependent variable. The inclusion of occupation, state ×year, and industry ×year fixed effects subsumes labor market competition, allowing us to attribute wage variation primarily to marginal productivity.

Columns (4)–(6) of Table 6 report the regression results. Greater value of the AI complementarity measure is associated with increased wages, validating the positive impact on marginal productivity. Specifically, in column (4), a one-percentile increase in AI net complementarity is associated with an increase of hourly wages by 0.34%, significant at the 1% level. The positive relationship remains robust across all specifications. Overall, the wage analysis suggests that working individuals, on average, experience positive financial gains from AI exposure due to the overall complementarity of AI technology to their human capital.

5.2. Performance monitoring: AI surveillance

Computerized workplace surveillance emerged in the 1980s (U.S. Congress, Office of Technology Assessment, 1987) and saw an unprecedented acceleration in 2020, driven by the shift to remote and hybrid work necessitated by COVID-19.[24] Technologies such as datafication, sensorization, and computer vision—backed by stronger cybersecurity infrastructure—have evolved from a stopgap solution during COVID-19 lockdowns into a scalable system for con-

[24]See, e.g., https://www.wsj.com/articles/youre-working-from-home-but-your-company-is-still-watching-you-11587202201?mod=Searchresults_pos20&page=1. Zuboff (2019) discusses how AI-powered platforms and digital tools create an environment of constant behavioral tracking and nudging, contributing to 24/7 responsiveness and the erosion of personal time.

tinuous remote supervision. This shift has reshaped organizational norms and management practices even after offices reopened. These technologies allow employers to capture increasingly precise, real-time signals of actual effort and performance with reduced noise. Within a principal-agent framework, such enhanced monitoring is expected to incentivize greater worker effort.

The 2020 COVID-19 shock provides a unique opportunity to assess how monitoring influences the length of the workday. Prior to the pandemic, remote work was often an endogenous choice; during the lockdowns, it became widespread. Remote and hybrid work persisted through the end of our sample period (2023), with the extent of remote work largely determined by job feasibility. To capture an ex ante effect, the sample is limited to occupations deemed remote-capable, defined by Dingel and Neiman (2020) as those whose pre-pandemic essential tasks did not require on-site presence. Within this subset of 65 occupations, we examine how the effectiveness of AI-based monitoring shapes worker effort.

Among individuals in occupations that can, ex ante, accommodate remote work, their reception to the AI surveillance technology shock in 2020 depends on the occupations' exposure to the new technology. Such an exposure could be constructed analogous to our main AI exposure measures. More specifically, we prompt the GPT model to assess how AI surveillance technologies enhance monitoring for each of the 65 occupations based on three applications of organizational control in the workplace: direction (restricting and recommending), evaluation(recording and rating), and discipline (replacing and rewarding), following Kellogg et al. (2020).[25] With the resulting surveillance exposure measure, AI_o^{SUR}, and the 2020 COVID-19 shock, we are able to conduct the following difference-in-differences estimation on observations indexed by individual (i), occupation (o), and year (t):

$$Y_{i,o,t} = \beta_1 \cdot AI_o^{SUR} \cdot Post_t + \beta_2 \cdot X_{i,t} + \alpha + \epsilon_{i,o,t}. \tag{5}$$

[25]Section B.4 of the Online Appendix describes the detailed procedures for measuring AI surveillance exposure of all occupations. Table OA.8 of the Online Appendix lists top occupations grouped by AI surveillance exposure.

The dependent variable is the number of weekly hours allocated to market work. The regression includes the same set of individual-level controls and fixed effects as in our baseline regressions. Since performance monitoring is a defining feature of a principal-agent setup and becomes moot in the absence of delegation, this hypothesis naturally lends itself to a placebo test: While AI surveillance technology is expected to elicit greater worker effort in equilibrium, the effect should be null for the self-employed.

Table 7 presents the weighted linear regression results for equation (6) for a sample of employees in ATUS. For individuals employed by a "principal" (i.e., employer), column (1) shows that a one-percentile increase in AI surveillance exposure is associated with a 0.044-hour increase in weekly work hours post-2020 (significant at the 5% level), which translates to 2.2 additional hours in a workweek for an interquartile variation.

[Insert Table 7 here.]

Columns (2)–(4) separately examine the effects of AI surveillance along three dimensions: direction, evaluation, and discipline. Both direction- and evaluation-based AI monitoring are associated with increased work hours, with coefficients significant at the 5% level and of similar magnitude to the general AI surveillance effect reported in column (1). In contrast, discipline-oriented AI technology shows no statistically significant relationship with work hours, although the coefficient remains positive and sizable (column (4)).

Monitoring is a feature of principal-agent relationships and employer-driven surveillance becomes largely irrelevant for the self-employed. Accordingly, the sub-sample of self-employed individuals (10.2% of ATUS respondents) serves as a comparison test, as reported in Table OA.8 of the Online Appendix. This sub-sample shows no statistically or economically significant relationship between AI surveillance exposure (or its components) and work hours.

5.3. Reservation utility: Employee welfare and market competition

5.3.1. Validating productivity surplus: Firm-level analysis

The AI-enabled labor productivity gain and prolonged work days discussed in earlier sections should manifest itself in stronger firm performance. This section validates the productivity rent using data from publicly traded firms. With a panel of firm (i) $\times$ year (t) level data from 2008 (the starting year of Glassdoor's coverage) to 2023, we estimate the following regression:

$$Y_{i,t} = \beta_1 \cdot AI^{EXP}_{i,t-1} + \beta_2 \cdot X_{i,t-1} + \alpha + \epsilon_{i,t}. \tag{6}$$

The dependent variable, $Y_{i,t}$ is firm-year level operating outcome. All regressions control for firm attributes, including sales in 2023 dollars (natural logarithm), Tobin's Q, market leverage, capital expenditure over he beginning-of-year assets, R&D expenditure over assets, asset tangibility (defined as net fixed assets over assets), firm fixed effects and year fixed effects. Standard errors are clustered at the firm level.

Table 8 presents the results. In columns (1)–(4), the dependent variable is ROA defined as operating income before depreciation over total assets, in percentage points. In columns (1) and (2), the key explanatory variables are workforce AI exposure ($AI^{EXP}_{i,t-1}$) and AI complementarity ($AI^{COMP}_{i,t-1}$), aggregated at the firm-year level using employment weights derived from the LinkedIn data. Firms enjoy higher ROA when their workers' job tasks become more exposed to or better complemented by AI, based on pre-existing employee composition and inferred from within-firm variation. Results are statistically significant at the 1% level. Columns (3) and (4) adopt predicted weekly work hours by lagged AI exposure ($\widehat{Work\_Hour}^{\text{EXP}}_{i,t-1}$) and AI complementarity ($\widehat{Work\_Hour}^{\text{COMP}}_{i,t-1}$) as key independent variables. Predicted work hours are generated based on the individual-level estimates from column (3) of Tables 3 and 6, and then aggregated to the occupation level using ATUS survey weights. Both columns confirm that when workers increase their work hours in response to AI exposure, firms benefit: a one-hour increase in average employee work hours is associated with a 13.5 basis point increase in ROA,

significant at the 5% level.

[Insert Table 8 here.]

Columns (5)–(8) examine an alternative firm outcome: labor productivity defined as sales over employment. The patterns are consistent with those for ROA, suggesting that technology-enabled labor productivity contribute to firms' profitability. Such gains could, in principle, be shared with workers through higher wages, improved benefits, or reduced long-term workloads. However, the extent to which these productivity gains are actually passed on to workers remains an open question—one we examine in subsequent analyses.

5.3.2. Employee welfare: Evidence from Glassdoor reviews

The relationship between technology-enabled productivity gains and workday length can also operate through the impact of these gains on workers' reservation utility, as agents re-optimize the allocation between work (and consequently, consumption) and leisure to adjust to a new welfare level determined within a competitive marketplace. At this new equilibrium, the effect of productivity gains—even when accompanied by higher compensation—on worker welfare remains a priori ambiguous, as factors such as self-motivation, fulfillment, and work-life balance all play critical roles in forming overall job satisfaction.

Employee reviews from Glassdoor (via Revelio, as detailed in Section 3) allow us to evaluate this relationship between worker satisfaction and AI exposure. Two metrics are most relevant: overall job satisfaction and Work-Life Balance (WLB) ratings. These metrics are measured at the occupation (o) $\times$ firm (i) $\times$ year (t) level, with both rating scales ranging from one (worst) to five (best). The sample covers both public and private firms. The main explanatory variable is lagged AI exposure at the occupation $\times$ year level. All specifications control for lagged employee review counts, the average seniority (ranging from 1-7), and remote work index (ranging from 0-1) of the occupation-by-firm cohort from Revelio, and fixed effects at the following level: occupation, firm $\times$ year. Standard errors are clustered by occupation.

Table 9 shows that greater AI exposure is associated with lower employee satisfaction (columns (1) and (2)), aligning with occupation-level evidence that AI exposure overall leads to extended work hours and decreased leisure time, despite the fact that wage increases (columns (3)) with productivity and work hours. Based on the coefficients in columns (1), inter-quartile increases in a firm's general AI exposure is associated with 2.6 basis point reductions in employees' overall satisfaction rating, respectively (average rating is 3.5). A qualitatively similar relationship is observed for the work-life balance (WLB) rating. An inter-quartile increase in AI exposure corresponds to a 2.45 basis point decrease in the WLB rating (average WLB rating is 3.4). Both effects are significant at the 1% level.

[Insert Table 9 here.]

Similar to Section 4.1, we also conduct an event study using a difference-in-differences framework around the release of ChatGPT. The test sample covers employee ratings from January 2022 to June 2023. Table OA.8 of Online Appendix presents the results, indicating that workers more exposed to AI report significantly lower satisfaction following the introduction of ChatGPT. We note that Berger et al. (2024) find negative but insignificant relation between employee overall/WLB rating and generative AI exposure at the firm level. The seeming difference could be reconciled by the different units of observations in their and our studies: the significantly negative relation in our setting is at the occupation level *within* firm-year.

We examine the sources of employees' dissatisfaction in two settings. The first test analyzes the textual content of Glassdoor reviews. Complaints, i.e., reviews that mention specific topics in a negative tone, are identified using the algorithm described in Section B.5 of the Online Appendix.[26] Column (4) of Table 9 shows that greater AI exposure is associated with more complaints about on-job surveillance, consistent with earlier findings that AI-driven surveillance technologies contribute to longer work hours. An inter-quartile increase

[26]Table OA.10 of the Online Appendix provides examples of complaints about surveillance and employment risk.

in AI exposure corresponds to a 0.95% increase in the number of complaints toward on-job surveillance based on the Poisson regressions for count data, equivalent to 8.6% of the sample mean (0.11 complaints per occupation-firm-year). On the other hand, AI exposure is not associated with increased complaints about employment risk (column (5)), aligning with evidence that AI exposure has limited effects on employment and separations during the sample period. This further supports our earlier finding that the extended work hours are unlikely driven by fear of job loss.

The second test connects leisure decomposition (demonstrated in Table 4) and personal marginal utility. Based on self-reported well-being data, Benjamin et al. (2025) estimate the marginal utility of 126 life aspects grouped into 15 domains. We link these life aspects with decomposed leisure activities using an LLM-based mapping procedure as described in Section B.6 of the Online Appendix. The mapping results reveal that leisure tends to enhance well-being in areas like "family well-being," "feelings," and "mental health," while potentially reducing utility in "status" driven activities. Applying respective utility parameters to each leisure activities based on the findings underlying Table 4, an interquartile increase in AI exposure is associated with approximately 5.5% loss of average leisure utility per week. This evidence suggests that leisure trade-offs meaningfully contribute to employee dissatisfaction in AI-exposed occupations.

Although the magnitude of the effects, a few basis-point changes in the ratings, may seem modest in isolation, the evidence clearly shows that employees have not reported an improved work experience, particularly in terms of work-life balance, when their jobs are more exposed to AI. If we treat these ratings as a proxy for worker welfare, the findings are disappointing given that these technologies are intended to make work more fulfilling and lives more enjoyable. Returning to our model, we interpret this result as evidence that, on average, worker reservation utility has not improved despite AI-enabled productivity gains.

5.3.3. Labor market competition: Worker bargaining power relative to firms

The extent to which workers' reservation utility rises with technology advancement depends on their relative bargaining power vis-à-vis the principal (employer). In a labor seller's market, workers are positioned to appropriate a greater share of the surplus from AI-enhanced productivity. An increase in reservation utility leads workers to work less relative to the level justified by increased productivity alone, analogous to an income effect. This hypothesis posits that the effect documented in Table 3 and Table 6 is expected to be weaker when workers have more bargaining power over firms.

Employment (i.e., labor buy side) concentration is a useful proxy for firms' monopsony power (e.g. Azar et al., 2020, 2022; Benmelech et al., 2022; Rinz, 2022). Following this literature, we measure employment concentration using the Herfindahl-Hirschman Index (HHI) at the state-occupation level. For each occupation o in state s and year t, HHI is calculated as the sum of squared employment shares of public firms in that occupation and state based on LinkedIn data. Higher HHI implies less intense labor market competition among firms. Another proxy for inter-firm competition is talent retention pressure (TRP), which reflects the challenges firms face in retaining skilled workers. Adapting the design developed by Chen et al. (2023) to our setting, we measure TRP with the job vacancy-to-employment ratio (V/E) at the state-occupation level, multiplied by each occupation's cognitive skill scores. The job vacancy data are from Burning Glass, employment data from OEWS, and cognitive skill scores from O*NET.[27] A higher TRP reflects greater retention pressure on firms, driven by competition for workers with advanced cognitive skills, who are generally scarce and possess abundant outside job opportunities. An indicator of high bargaining power of workers relative to firms, *I(Worker Power vs. Firm)*, is set to one if labor market competition among firms is high (i.e, employment HHI is in the bottom quartile or TRP is in the top quartile), and zero otherwise.

Columns (1) and (2) of Table 10 present the heterogeneous effects of AI net complemen-

[27]The job vacancy data is from 2010 to 2018 and we identify unique job vacancies following Jiang et al. (2025). The cognitive skill score of each occupation, ranging from one to five, is computed as the average importance of six cognitive skills using the O*NET 23.0 database (Acemoglu and David, 2011; Chen et al., 2023).

tarity on work hours, conditional on workers' bargaining power relative to firms. Both proxies of employee's bargaining power are associated with a significantly smaller increase in work hours as AI complementarity rises. Specifically, with a an interquartile increase in AI complementarity, a worker in the bottom quartile of employment concentration exhibits a smaller increase in work hours by 0.75 (column (1), significant at the 10% level). Similarly, a worker in the top quartile of TRP shows a smaller increase in work hours by 3.3 (column (2), significant at the 5% level). These results indicate that greater worker bargaining power attenuates the positive association between AI complementarity and work hours. Replacing $AI_{o,t-1}^{COMP}$ with $AI_{o,t-1}^{EXP}$ yields qualitatively similar findings (see Table OA.9 of the Online Appendix.)

[Insert Table 10 here.]

5.3.4. Product market competition: Firm bargaining power over consumers

Parallel to labor market competition, a firm's product market power determines how productive surplus is split between firms and the consumer of their products or services. The more pricing power of firms relative to consumers, potentially more surplus may eventually accrues to labor as there is more to split among parties on the production side, which is expected to mitigate the impact of AI net complementarity on workday via an income effect.

We adopt two measures for firms' product market power, both provided by Hoberg and Phillips (2016) based on public firms. One is firm-level product similarity that assesses how closely a firm's product descriptions in its 10-K filings match those of industry peers. The other is the firm-level HHI, defined as the sum of the squared market shares of firms in the same 10-K text-based industry using Compustat sales data.[28] Higher HHI and lower product similarity suggest greater firm bargaining power over consumers, and potential surplus sharing with workers. The indicator for high pricing power of firms in an industry relative to consumers, *I(Firm Power vs. Consumer)*, is set to one if the lagged product similarity is in the bottom

[28]Note that ATUS lacks firm identifiers and reports only Census industry codes, and Compustat firms are linked to the NAICS classification. To link the two, we first calculate the sales-weighted product market competition proxies at the NAICS 3-digit industry level, and then match them to the corresponding Census industry code "trdtind1" used in ATUS using the crosswalk provided by BLS at https://www2.census.gov/programs-surveys/demo/guidance/industry-occupation/census-2012-final-code-list.xls.

quartile or if the product HHI is in the top quartile, and zero otherwise. Each of these two indicators are interacted with occupational AI net complementarity.

Columns (3) and (4) of Table 10 present weighted linear regression estimates of the heterogeneous effect of AI complementarity on workdays, interacted with product market power. A higher value of *I(Firm Power vs. Consumer)* weakens the positive relationship between an interquartile increase in AI net complementarity and weekly work hours by 0.6 in column (3) (statistically insignificant, t=-1.50) and by 1.55 in column (4) (significant at the 1% level). These results offer some evidence that when firms hold greater bargaining power over consumers, the impact of AI technologies on work hours is partially mitigated. However, this trickle-down benefit to workers appears weaker than the effect of workers' direct bargaining power over firms.

5.3.5. Impact of employment risk

Though out of our sample period, the widespread concern about the potential mass layoff driven by AI since 2024 warrants a separate discussion. First, we note that our default sample explicitly excluded the technology sector, the main destination of mass layoffs since 2024.[29] Moreover, the unemployment rate released by BLS for June 2025 stood at 4.1% and has remained in a narrow range of 4.0% to 4.2% since May 2024. This suggests that, overall, the labor market has not experienced the AI-triggered employment shock economy wide.

We nevertheless conduct a formal test on employment risk building on LinkedIn employment history data. We define *Hire* as an indicator that equals one if an employee is hired by a firm in a given year, while *Separation* is an indicator that equals one if an employee leaves the firm in a given year. Further, a departure is considered voluntary if the employee joins a firm with a larger (more than 25%) employment or moves to a role with higher seniority, following Jiang et al. (2025). Next, we aggregate the number of hires and separations to the occupation (o) $\times$ state (s) $\times$ year (t) level.

[29] According to layoffs.fyi, a website that tracks job cuts in the tech sector, from January through July 2025 over 70,000 tech workers from over 100 companies, including Amazon, Meta, Microsoft, Intel, etc., have been laid off.

Table 11 presents regression results where the dependent variable is the natural logarithm of hires and separations. The variables of interest are the lagged AI exposure. All columns incorporate lagged controls at the occupation-by-state cohort level, including LinkedIn employment count (natural logarithm), average seniority and remote work index from Revelio, and fixed effects at the occupation and the state × year level. The regression is weighted by the underlying LinkedIn employment count at the occupation-by-state cohort level, and standard errors are double clustered by occupation and state.

[Insert Table 11 here.]

Columns (1) and (5) show that AI exposure and AI net complementarity are positively but insignificantly related to new hires, while columns (2) and (6) report similar patterns for separations. These muted effects are consistent with the employment trends in Table OA.7 of the Online Appendix and findings in the recent literature(e.g., Acemoglu et al., 2022; Hampole et al., 2025). Breaking down separations, column (3) shows that an interquartile increase in AI exposure raises voluntary separations by 3.65%, significant at the 5% level. This effect is more pronounced for AI net complementarity, approximately doubling in magnitude and significant at the 1% level (column (7)). That is, AI exposure, particularly when complementary, increases external opportunities. In contrast, no significant relation is found for involuntary separations (columns (4) and (8)), consistent with the previously reported evidence that AI adoption had not triggered disproportionate layoffs as of 2023. Overall, these results suggest that longer work hours are unlikely driven by fear of employment risk or task consolidation upon downsizing.

6. Conclusion

The extensive individual-level time diary data (ATUS) collected over the past two decades offers a unique setting to examine the nuanced relationship between occupational AI exposure and workers' time allocation. Our analysis reveals a consistent pattern: workers in occupations

with higher AI exposure end up working longer hours and enjoying less leisure time. This effect is particularly pronounced in contexts where AI significantly enhances marginal productivity and monitoring efficiency. It is further amplified in competitive labor and product markets, where workers' limited bargaining power fails to keep up with productivity gains, with rents often accruing to firms or consumers.

Historically, technological advancements like the Industrial Revolution and automation initially increased work hours as productivity demands rose and labor shifted to factory-based systems. Over time, however, productivity gains and social reforms reduced work hours, especially in developed economies, enabling improved work-life balance. Such a historical trend has contributed to the expectation for AI technologies. Our findings challenge the prevailing goal and assumption that technology progress improves lives including alleviating human labor burdens. Instead, they uncover a paradox where AI-driven productivity gains and enhanced monitoring efficiency extend workdays, especially in contexts with limited opportunities for workers to share in the benefits. To achieve a world where humans work less and enjoy greater well-being, deliberate policy interventions, equitable distribution of productivity gains, and cultural shifts prioritizing leisure and quality of life are essential. By shedding light on AI's impact on work-life dynamics from a principal-agent framework, this study contributes to the broader discussion on the socio-economic consequences of emerging technologies.

References

Acemoglu, Daron, 1998, Why do new technologies complement skills? directed technical change and wage inequality, *The quarterly journal of economics* 113, 1055–1089.

Acemoglu, Daron, David Autor, Jonathon Hazell, and Pascual Restrepo, 2022, Artificial intelligence and jobs: Evidence from online vacancies, *Journal of Labor Economics* 40, S293–S340.

Acemoglu, Daron, and Autor David, 2011, Skills, tasks and technologies: Implications for employment and earnings, *Handbook of labor economics* 4, 1044–1171.

Acemoglu, Daron, and Pascual Restrepo, 2019, Automation and new tasks: How technology

displaces and reinstates labor, *Journal of economic perspectives* 33, 3–30.

Acemoglu, Daron, and Pascual Restrepo, 2022, Tasks, automation, and the rise in us wage inequality, *Econometrica* 90, 1973–2016.

Aguiar, Mark, Mark Bils, Kerwin Kofi Charles, and Erik Hurst, 2021, Leisure luxuries and the labor supply of young men, *Journal of Political Economy* 129, 337–382.

Aguiar, Mark, and Erik Hurst, 2007, Measuring trends in leisure: The allocation of time over five decades, *The quarterly journal of economics* 122, 969–1006.

Aguiar, Mark, Erik Hurst, and Loukas Karabarbounis, 2013, Time use during the great recession, *American Economic Review* 103, 1664–1696.

Alon, Titan, Matthias Doepke, Jane Olmstead-Rumsey, and Michèle Tertilt, 2020, This time it's different: the role of women's employment in a pandemic recession, Technical report, National Bureau of Economic Research.

Armour, John, Richard Parnham, and Mari Sako, 2022, Augmented lawyering, *University of Illinois Law Review* .

Autor, David, and David Dorn, 2009, This job is "getting old": measuring changes in job opportunities using occupational age structure, *American Economic Review* 99, 45–51.

Autor, David, and David Dorn, 2013, The growth of low-skill service jobs and the polarization of the us labor market, *American Economic Review* 103, 1553–1597.

Autor, David H, 2015, Why are there still so many jobs? the history and future of workplace automation, *Journal of Economic Perspectives* 29, 3–30.

Azar, José, Ioana Marinescu, and Marshall Steinbaum, 2022, Labor market concentration, *Journal of Human Resources* 57, S167–S199.

Azar, José, Ioana Marinescu, Marshall Steinbaum, and Bledi Taska, 2020, Concentration in us labor markets: Evidence from online vacancy data, *Labour Economics* 66, 101886.

Babina, Tania, Anastassia Fedyk, Alex He, and James Hodson, 2024, Artificial intelligence, firm growth, and product innovation, *Journal of Financial Economics* 151, 103745.

Bandiera, Oriana, Andrea Prat, Stephen Hansen, and Raffaella Sadun, 2020, Ceo behavior and firm performance, *Journal of Political Economy* 128, 1325–1369.

Becker, Gary S, 1965, A theory of the allocation of time, *The economic journal* 75, 493–517.

Ben-Rephael, Azi, Bruce I Carlin, Zhi Da, and Ryan D Israelsen, 2025, Uncovering the hidden

effort problem, *The Journal of Finance* .

Benjamin, Daniel J, Kristen B Cooper, Ori Heffetz, Miles S Kimball, and Tushar Kundu, 2025, What do people want?, Technical report, National Bureau of Economic Research.

Benmelech, Efraim, Nittai K Bergman, and Hyunseob Kim, 2022, Strong employers and weak employees: How does employer concentration affect wages?, *Journal of Human Resources* 57, S200–S250.

Berger, Philip G, Wei Cai, Lin Qiu, and Cindy Xinyi Shen, 2024, Employer and employee responses to generative ai: Early evidence, *Available at SSRN 4874061* .

Blundell, Richard, and Thomas MaCurdy, 1999, Labor supply: A review of alternative approaches, in Orley Ashenfelter, and David Card, eds., *Handbook of Labor Economics*, volume 3A, 1559–1695 Elsevier.

Boerma, Job, and Loukas Karabarbounis, 2021, Inferring inequality with home production, *Econometrica* 89, 2517–2556.

Briot, Jean-Pierre, 2021, From artificial neural networks to deep learning for music generation: history, concepts and trends, *Neural Computing and Applications* 33, 39–65.

Briot, Jean-Pierre, Gaëtan Hadjeres, and François-David Pachet, 2017, Deep learning techniques for music generation–a survey, *arXiv preprint arXiv:1709.01620* .

Brogaard, Jonathan, Matthew C Ringgenberg, and Dominik Roesch, 2024, Does floor trading matter?, *The Journal of Finance* .

Cao, Sean, Wei Jiang, Junbo Wang, and Baozhong Yang, 2024, From man vs. machine to man + machine: The art and ai of stock analyses, *Journal of Financial Economics* 160.

Casey, Anthony J, and Anthony Niblett, 2016, Self-driving laws, *University of Toronto Law Journal* 66, 429–442.

Chen, AJ, Miao Ben Zhang, and Zhao Zhang, 2023, Talent market competition and firm growth, *Available at SSRN 4597388* .

Cheng, Xin, Evgeny Lyandres, Kaiguo Zhou, and Tong Zhou, 2024, Labor-replacing automation and finance, *Management Science* .

Commerford, Benjamin P, Sean A Dennis, Jennifer R Joe, and Jenny W Ulla, 2022, Man versus machine: Complex estimates and auditor reliance on artificial intelligence, *Journal of Accounting Research* 60, 171–201.

Cook, Cody, Rebecca Diamond, Jonathan V Hall, John A List, and Paul Oyer, 2021, The gender earnings gap in the gig economy: Evidence from over a million rideshare drivers, *The Review of Economic Studies* 88, 2210–2238.

Cortés, Patricia, Ying Feng, Nicolás Guida-Johnson, and Jessica Pan, 2024, Automation and gender: Implications for occupational segregation and the gender skill gap, Technical report, National Bureau of Economic Research.

David, Autor, and Wasserman Melanie, 2013, Wayward sons: The emerging gender gap in labor markets and education, *Third Way Report* 20013.

de Kok, Ties, 2025, Chatgpt for textual analysis? how to use generative llms in accounting research, *Management Science* .

de la Rosa, Pedro A, Richard G Cowden, Joseph A Bulbulia, Chris G Sibley, and Tyler J VanderWeele, 2024, Effects of screen-based leisure time on 24 subsequent health and well-being outcomes: A longitudinal outcome-wide analysis, *International Journal of Behavioral Medicine* 1–20.

Del Boca, Daniela, and Christopher Flinn, 2012, Endogenous household interaction, *Journal of Econometrics* 166, 49–65.

Dingel, Jonathan I., and Brent Neiman, 2020, How many jobs can be done at home?, *Journal of Public Economics* 189.

Doepke, Matthias, Anne Hannusch, Fabian Kindermann, and Michèle Tertilt, 2023, The economics of fertility: A new era, in *Handbook of the Economics of the Family*, volume 1, 151–254 Elsevier.

Dorn, David, 2009, *Essays on inequality, spatial interaction, and the demand for skills*, Ph.D. thesis, Verlag nicht ermittelbar.

Edmans, Alex, 2011, Does the stock market fully value intangibles? employee satisfaction and equity prices, *Journal of Financial economics* 101, 621–640.

Eisfeldt, Andrea L, Gregor Schubert, Miao Ben Zhang, and Bledi Taska, 2023, Generative ai and firm values, *Available at SSRN 4436627* .

Fedyk, Anastassia, and James Hodson, 2023, Trading on talent: Human capital and firm performance, *Review of Finance* 27, 1659–1698.

Fehr, Ernst, and Lorenz Goette, 2007, Do workers work more if wages are high? evidence

from a randomized field experiment, *American Economic Review* 97, 298–317.

Felten, Edward W, Manav Raj, and Robert Seamans, 2018, A method to link advances in artificial intelligence to occupational abilities, in *AEA Papers and Proceedings*, volume 108, 54–57, American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203.

Felten, Edward W, Manav Raj, and Robert Seamans, 2019, The occupational impact of artificial intelligence: Labor, skills, and polarization, *NYU Stern School of Business* .

Firpo, Sergio, Nicole M Fortin, and Thomas Lemieux, 2011, Occupational tasks and changes in the wage structure, IZA Discussion Paper.

Gornall, Will, Oleg R Gredil, Sabrina T Howell, Xing Liu, and Jason Sockin, 2024, Do employees cheer for private equity? the heterogeneous effects of buyouts on job quality, *Management Science* .

Graff Zivin, Joshua, and Matthew Neidell, 2014, Temperature and the allocation of time: Implications for climate change, *Journal of Labor Economics* 32, 1–26.

Green, T Clifton, Ruoyan Huang, Quan Wen, and Dexin Zhou, 2019, Crowdsourced employer reviews and stock returns, *Journal of Financial Economics* 134, 236–251.

Gu, Shihao, Bryan Kelly, and Dacheng Xiu, 2020, Empirical asset pricing via machine learning, *The Review of Financial Studies* 33, 2223–2273.

Hampole, Menaka, Dimitris Papanikolaou, Lawrence D.W. Schmidt, and Bryan Seegmiller, 2025, Artificial intelligence and the labor market, Working Paper 33509, National Bureau of Economic Research.

Hartley, Jonathan, Filip Jolevski, Vitor Melo, and Brendan Moore, 2024, The labor market effects of generative artificial intelligence, *Available at SSRN* .

Hoberg, Gerard, and Gordon Phillips, 2016, Text-based network industries and endogenous product differentiation, *Journal of political economy* 124, 1423–1465.

Holmstrom, Bengt, and Paul Milgrom, 1987, Aggregation and linearity in the provision of intertemporal incentives, *Econometrica* 55, 303–328.

Hui, Xiang, Oren Reshef, and Luofeng Zhou, 2024, The short-term effects of generative artificial intelligence on employment: Evidence from an online labor market, *Organization Science* 35, 1977–1989.

Jevons, William Stanley, 1865, *The Coal Question; An Inquiry Concerning the Progress of the*

Nation, and the Probable Exhaustion of Our Coal Mines Macmillan & Co., London, United Kingdom.

Jiang, Wei, Yuehua Tang, Rachel J Xiao, and Vincent Yao, 2025, Surviving the fintech disruption, *Journal of Financial Economics* 171, 104071.

Kellogg, Katherine C, Melissa A Valentine, and Angele Christin, 2020, Algorithms at work: The new contested terrain of control, *Academy of management annals* 14, 366–410.

Kogan, Leonid, Dimitris Papanikolaou, Lawrence DW Schmidt, and Bryan Seegmiller, 2023, Technology and labor displacement: Evidence from linking patents with worker-level data, Technical report, National Bureau of Economic Research.

Kogan, Leonid, Dimitris Papanikolaou, Amit Seru, and Noah Stoffman, 2017, Technological innovation, resource allocation, and growth, *The Quarterly Journal of Economics* 132, 665–712.

Krueger, Alan B, and Andreas Mueller, 2010, Job search and unemployment insurance: New evidence from time use data, *Journal of Public Economics* 94, 298–307.

Liu, Yuanyang, Chuanren Liu, and Tingliang Huang, 2023, Does ai reduce inequality? a study with a new occupational ai exposure measure, SSRN Working Paper 4515629.

Lopez-Lira, Alejandro, and Yuehua Tang, 2023, Can chatgpt forecast stock price movements? return predictability and large language models, *arXiv preprint arXiv:2304.07619* .

MaCurdy, Thomas E., 1981, An empirical model of labor supply in a life-cycle setting, *Journal of Political Economy* 89, 1059–1085.

Pairolero, Nicholas A, Alexander V Giczy, Gerard Torres, Tisa Islam Erana, Mark A Finlayson, and Andrew A Toole, 2025, The artificial intelligence patent dataset (aipd) 2023 update, *The Journal of Technology Transfer* 1–24.

Pennington, Jeffrey, Richard Socher, and Christopher D Manning, 2014, Glove: Global vectors for word representation, in *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, 1532–1543.

Rinz, Kevin, 2022, Labor market concentration, earnings, and inequality, *Journal of Human Resources* 57, S251–S283.

Surden, Harry, 2019, Artificial intelligence and law: An overview, *Georgia State University Law Review* 35.

Tuzel, Selale, and Miao Ben Zhang, 2021, Economic stimulus at the expense of routine-task jobs, *The Journal of Finance* 76, 3347–3399.

U.S. Congress, Office of Technology Assessment, 1987, The electronic supervisor: New technology, new tensions, Technical Report OTA-CIT-333, U.S. Government Printing Office, Washington, DC.

Wang, Weiguang, Guodong Gao, and Ritu Agarwal, 2024, Friend or foe? teaming between artificial intelligence and workers with variation in experience, *Management Science* 70, 5753–5775.

Webb, Michael, 2019, The impact of artificial intelligence on the labor market, SSRN Working Paper 3482150.

Yang, Chih-Hai, 2022, How artificial intelligence technology affects productivity and employment: firm-level evidence from taiwan, *Research Policy* 51, 104536.

Zhang, Miao Ben, 2019, Labor-technology substitution: Implications for asset pricing, *The Journal of Finance* 74, 1793–1839.

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

Figure 1. Most Frequent Keywords in AI Patents

A: AI Patents 2000 - 2009

B: AI Patents 2010 - 2019

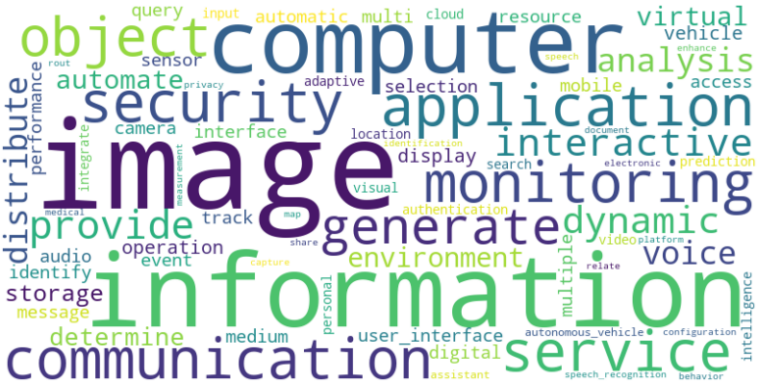

C: AI Patents 2020 - 2023

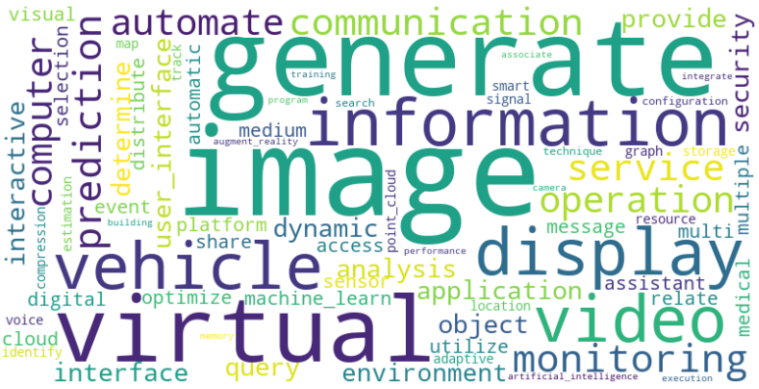

This figure presents word clouds of the top keywords appearing in AI patent titles during three periods: 2000–2009 (Panel A), 2010–2019 (Panel B), and 2020–2023 (Panel C). Font size is proportional to keyword frequency. Each panel displays the 100 most frequent keywords, excluding generic terms such as "system," "method," and "device."

Figure 2. Examples of High-Impact AI Patents Over Time

The timeline shows examples of high-impact AI patents and their business applications from 2000 to 2023. High-impact AI patents are defined as those in the top 1% every year based on their adjusted forward citations. Following Kogan et al. (2017), a patent's adjusted forward citations are calculated as its raw citation count divided by the average citation count of AI patents granted in the same year-quarter and CPC subclass.

Figure 3. AI Exposure Over Time

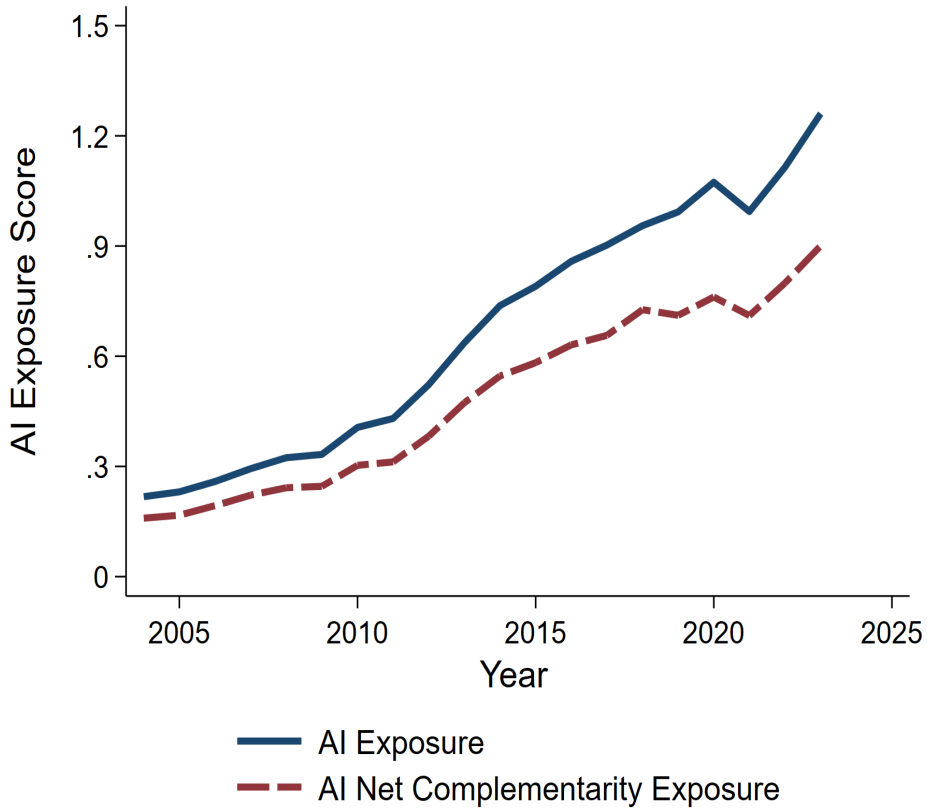

The figure plots the average occupational AI exposure of ATUS respondents over time. The average is calculated using ATUS survey weights. Two AI exposure measures are constructed by the authors using job task descriptions and AI patents published in the five years ending in a given year from 2000 to 2023: (i) average AI exposure based on the overlap of job tasks of occupations and AI patents (blue line) and (ii) average AI net complementarity exposure (red dotted line). Section B.2 describes the variable construction.

Figure 4. AI Exposure and Workday

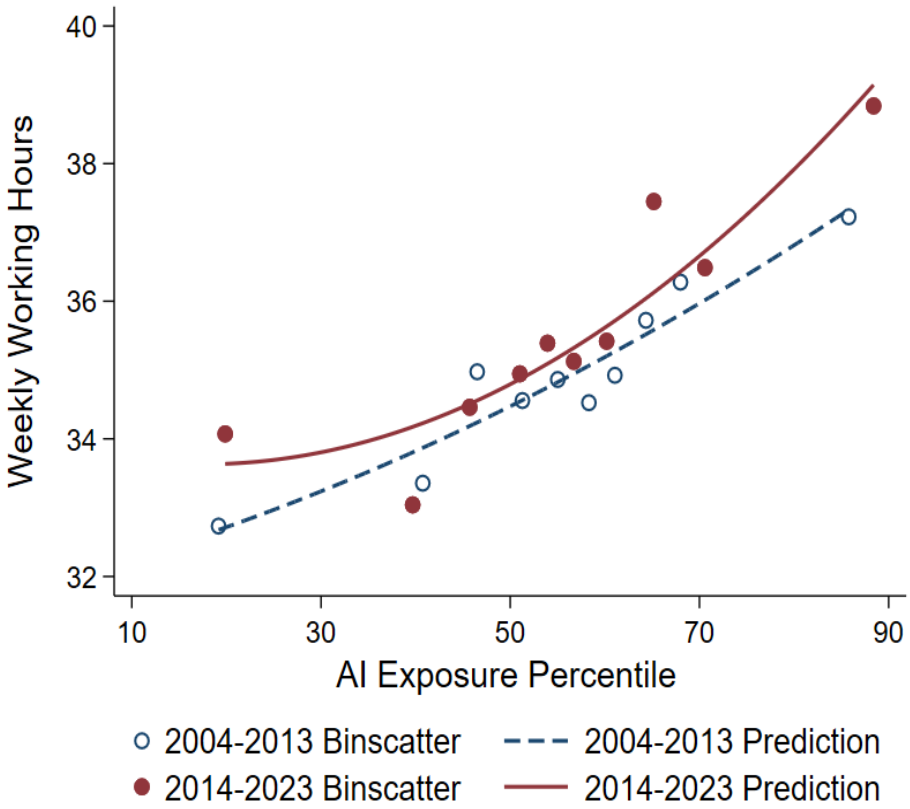

The figure plots the average weekly work hours over occupational AI exposure. The time allocation variables are derived from the American Time Use Survey (ATUS) for the periods 2004–2013 and 2014–2023, weighted using ATUS sampling weights. Blue scatters and the blue dotted line represent data from 2004–2013, while red scatters and the red line correspond to 2014–2023. The scatters depict binned averages, and the lines show fitted values from quadratic regressions, both adjusted for occupation group effects. Occupational AI exposure is constructed by the authors using job task descriptions and AI patents published in a five-year rolling window. The raw score for AI exposure is transformed into percentile ranks by year following the literature (e.g., Autor and Dorn, 2013; Webb, 2019).

Table 1: Summary Statistics

Panel A describes the individual-level variables in the ATUS sample from 2004 to 2023 and the means are calculated using ATUS sample weights following Aguiar and Hurst (2007). The occupations are uniquely identified by "*occ*1990*dd*" codes from Dorn (2009). The time spent on activities is from the ATUS, expressed in hours per week. An individual's total time endowment, after subtracting off 49 hours for biological eating, sleeping, and personal care needs (ESP), is 119 hours per week. *Market work* includes time spent on main jobs, overtime work, and ancillary work activities. *Leisure* includes entertainment like recreational computing and video games, hobbies and leisure components of ESP. *Home production* includes household chores, grocery shopping, caring for other adults, etc. *Education* refers to one's own education like attending courses. Civic includes going to church, volunteering, etc. *Job search* includes submitting resumes and conducting job interviews. *Hourly wages* are in 2023 dollars. The time-varying exposure measures at the occupation level, including AI exposure (AI^{EXP}) and AI net complementarity exposure (AI^{COMP}), are constructed by the authors and transformed into percentile ranks by year, as described in Section 3.6. $GenAI^{EXP}$ is generative AI exposure measure at the occupation level constructed following Eisfeldt et al. (2023). General AI surveillance exposure, AI^{SUR}, and its decomposed measures along the dimensions of direction, evaluation, and discipline ($AI^{SUR} - Direction$, $AI^{SUR} - Evaluation$, and $AI^{SUR} - Discipline$), are described in Section B.4 of Online Appendix. Panel B summarizes employee ratings, annual salaries, seniority, and remote work index at the firm-occupation level from 2008 to June 2023. Employees' ratings on overall satisfaction and work-life-balance (WLB) are from Glassdoor; annual salaries in 2023 dollars, seniority level (from 1 to 7), and remote potentials (ranging from 0 to 1) are derived by Revelio.

Panel A: Occupation Exposure, Time Allocation and Wages at the Individual Level

VARIABLES	N	Mean	Std	P25	P50	P75
	(1)	(2)	(3)	(4)	(5)	(6)
Weekly Hours						
Market work	124,385	35.21	30.42	0	45.03	59.50
Leisure	124,385	55.28	27.23	35.93	50.52	72.33
Education	124,385	1.20	5.95	0	0	0
Civic activities	124,385	1.60	5.88	0	0	0
Own medical care	124,385	0.37	2.10	0	0	0
Job Search	124,385	0.08	0.87	0	0	0
Home production	124,385	15.85	17.43	2.33	10.03	23.33
Child care	124,385	4.55	11.26	0	0	2
Hourly wages ($)	104,779	28.13	17.78	15.43	22.64	35.38
I(Female)	124,385	0.48	0.50	0	0	1
I(Married)	124,385	0.56	0.50	0	1	1
No. Children	124,385	0.80	1.12	0	0	1
Age	124,385	40.60	12.82	30.00	41.00	51.00
Indicator for Educational Attachment						
I(Less than high school)	124,385	0.09	0.28	0	0	0
I(High school)	124,385	0.29	0.45	0	0	1
I(Some college education)	124,385	0.23	0.42	0	0	0
I(Bachelor's)	124,385	0.27	0.44	0	0	1
I(Master's and above)	124,385	0.13	0.34	0	0	0
AI^{EXP}- score	124,385	0.66	0.39	0.32	0.61	0.96
AI^{COMP}- score	124,385	0.48	0.38	0.19	0.36	0.75
$GenAI^{EXP}$- score	8,185	0.38	0	0	0	1
AI^{SUR} - score	11,120	0.52	0.19	0	1	1
$AI^{SUR} - Direction$	11,120	0.49	0.18	0	0	1
$AI^{SUR} - Evaluation$	11,120	0.51	0.18	0	1	1
$AI^{SUR} - Discipline$	11,120	0.53	0.18	0	1	1

Panel B: Summary Statistics at the Occupation × Firm Level						
VARIABLES	N	Mean	Std	P25	P50	P75
	(1)	(2)	(3)	(4)	(5)	(6)
Rating - overall	436,858	3.50	1.09	3	4	4
Rating - WLB	408,239	3.40	1.15	3	4	4
No.Complaints - Surveillance	433,998	0.11	0.55	0	0	0
No.Complaints - Employment Risk	433,998	0.03	0.23	0	0	0
Annual salary ($000)	436,858	87.71	44.24	55	83	110
No.Reviews	436,858	4.75	9.69	1	2	4
Seniority	436,858	2.56	1.01	2	3	3
Remote Work Index	436,858	0.46	0.19	0	0	1

Table 2: Event Study: Introduction of ChatGPT

The table reports the weighted linear regressions that examine the effect of occupational exposure to generative AI on work and leisure based on individual responses to the ATUS survey from 2022 to 2023 using the ATUS sample weights (Aguiar et al., 2021). The occupation classification is the same as in Table 1. The dependent variable is weekly hours spent on market work in Panel A and leisure in Panel B. In each panel, column (1) presents the results for the full sample. Columns (2)–(3) present the results for subsamples defined using generative AI complementarity exposure at the occupation level, developed following Kogan et al. (2023). Columns (4)–(5) present the results for subsamples defined using the state-level Google search trend of ChatGPT from November 30 to December 31, 2022. Columns (6)–(7) present results for subsamples divided by their overwork status, where the "already overworking" workers are from occupations whose usual work hours per week were no fewer than 45 hours in 2021. The main explanatory variable, $GenAI^{EXP}$, is generative AI exposure measure at the occupation level from Eisfeldt et al. (2023) and transformed to percentile ranks following the literature (e.g., Autor and Dorn, 2013; Webb, 2019). $POST$ dummy equals one for the year 2023. All specifications include individual-level controls including age, the number of children, and a series of indicator variables for gender, educational attainment, marital status, and race, and fixed effects at the following levels: occupation, state × year, industry × year, year-month, and day-of-week. Standard errors are double clustered by occupation and state. Asterisks denote significance levels (***=1%, **=5%, *=10%).

Panel A: Work							
Dep Var	Weekly Work Hours$_{i,o,t}$						
		$GenAI_o^{COMP}$		State-level Google Search of ChatGPT$_s$		Already Overworking$_o$	
Sample	Full Sample	Top 25%	Bottom 75%	Top 25%	Bottom 75%	Yes	No
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
$GenAI_o^{EXP} \times POST_t$	0.075**	0.177**	0.090**	0.118*	0.078***	-0.036	0.078***
	(2.57)	(2.08)	(2.27)	(1.92)	(2.72)	(-0.21)	(3.01)
Individual Characteristics	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Occupation FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
State × Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry × Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year-Month FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Day-of-Week FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	8,148	3,512	4,628	1,946	6,126	925	7,141
R^2	0.343	0.471	0.318	0.408	0.368	0.491	0.343
Adjusted R^2	0.301	0.421	0.253	0.298	0.318	0.323	0.301

Panel B: Leisure							
Dep Var	Weekly Leisure Hours$_{i,o,t}$						
		$GenAI_o^{COMP}$		State-level Google Search of ChatGPT$_s$		Already Overworking$_o$	
Sample	Full Sample	Top 25%	Bottom 75%	Top 25%	Bottom 75%	Yes	No
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
$GenAI_o^{EXP} \times POST_t$	-0.077*** (-2.73)	-0.152** (-2.15)	-0.100** (-2.53)	-0.099* (-1.93)	-0.091*** (-2.87)	0.008 (0.06)	-0.092*** (-3.21)
Individual Characteristics	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Occupation FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
State × Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry × Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year-Month FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Day-of-Week FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	8,148	3,512	4,628	1,946	6,126	925	7,141
R^2	0.291	0.365	0.294	0.364	0.313	0.483	0.293
Adjusted R^2	0.247	0.305	0.226	0.245	0.258	0.313	0.247

Table 3: AI Exposure and Workday

The table reports weighted linear regression results that examine the impact of AI exposure on work and leisure based on individual responses to the ATUS survey from 2004–2023 using ATUS sample weights (Aguiar et al., 2021). The occupation classification is the same as in Table 1. The dependent variable is weekly hours spent on market work in columns (1)–(3) and leisure in columns (4)–(6). The main explanatory variable, AI^{EXP}, is AI exposure measure in percentile rank at the occupation -year level, calculated from AI-related patents granted over five years ending in the current year (detailed description in Section 3.6). All specifications control for individual characteristics, including age, the number of children, and a series of indicator variables for gender, educational attainment, marital status, and race, and fixed effects at the following levels: state × year, industry × year, year-month and day-of-week. Columns (2) and (5) additionally control for other occupational exposure measures, including robot exposure (Webb, 2019), routine task index (RTI) (Autor and Dorn, 2013), and offshorability exposure (Firpo et al., 2011; Autor and Dorn, 2013), all in percentile ranks. Columns (3) and (6) include occupation fixed effects, which subsume occupation-level controls. Standard errors are double clustered by occupation and state. Asterisks denote significance levels (***=1%, **=5%, *=10%).

Dep Var	Weekly Hours$_{i,o,t}$					
	Work			Leisure		
	(1)	(2)	(3)	(4)	(5)	(6)
$AI^{EXP}_{o,t-1}$	0.045***	0.036***	0.023**	-0.031***	-0.025***	-0.019***
	(3.55)	(3.20)	(2.58)	(-3.15)	(-2.92)	(-3.17)
Individual characteristics	Yes	Yes	Yes	Yes	Yes	Yes
Other Occupation Exposure	No	Yes	No	No	Yes	No
Occupation FE	No	No	Yes	No	No	Yes
State × Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Industry × Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Year-Month FE	Yes	Yes	Yes	Yes	Yes	Yes
Day-of-Week FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	124,059	124,059	124,059	124,059	124,059	124,059
R^2	0.274	0.276	0.285	0.235	0.236	0.243
Adjusted R^2	0.261	0.264	0.270	0.221	0.223	0.228

Table 4: Decomposed Leisure Activities

The table reports the weighted linear regressions that estimate the effect of occupational AI exposure on leisure activities at the individual level. The ATUS survey sample and occupation classification are the same as in Table 3. The dependent variable, weekly hours spent on leisure activities, is categorized into screen-based leisure activities (recreational computer use, gaming, and watching TV) in column (1), and non-screen leisure activities in column (2). Column (3)–(6) further decompose the non-screen leisure activities subdivided into four categories: recreation (relaxing, listening to music, traveling, etc.), socializing, leisure aspects of eating, sleeping, and personal care (ESP), and others (hobbies, reading, and sports). The main explanatory variable is AI exposure in percentile rank at the occupation-year level, calculated from AI-related patents granted over the five years ending in the current year (detailed description in Section 3.6). We additionally control for individual characteristics, including age, the number of children, and a series of indicator variables for gender, educational attainment, marital status, and race, and fixed effects at the following levels: occupation, state × year, industry × year, year-month, and day-of-week. Standard errors are double clustered by occupation and state. Asterisks denote the significance levels (***=1%, **=5%, *=10%).

Dep Var	Weekly Leisure Hours$_{i,o,t}$					
	Screen-Based	Non-Screen	Non-Screen			
			Recreation	Socializing	ESP	Other
	(1)	(2)	(3)	(4)	(5)	(6)
$AI^{EXP}_{o,t-1}$	-0.001	-0.018***	-0.005**	-0.004	-0.010*	0.001
	(-0.24)	(-3.38)	(-2.08)	(-0.93)	(-1.85)	(0.25)
Individual characteristics	Yes	Yes	Yes	Yes	Yes	Yes
Occupational FE	Yes	Yes	Yes	Yes	Yes	Yes
State × Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Industry × Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Year-Month FE	Yes	Yes	Yes	Yes	Yes	Yes
Day-of-Week FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	124,059	124,059	124,059	124,059	124,059	124,059
R^2	0.132	0.153	0.051	0.076	0.135	0.078
Adjusted R^2	0.114	0.136	0.031	0.058	0.118	0.059

Table 5: AI Exposure and Workday: Within-Occupation Dispersion

The table reports the weighted linear regression results that examine the effect of AI on the within-occupation variation of workdays at the occupation × year level derived from the ATUS data from 2004-2023. The regression is weighted by ATUS sample weights at the occupation-year level. The occupation classification is the same as in Table 3. The dependent variable is the differences in weekly work hours between females and males in column (1)-(2) and between young and old workers in columns (3)–(4). Young workers are defined as those whose age is in the bottom quartile within the occupation in a given year. The main explanatory variable is AI exposure in percentile rank at the occupation-year level, calculated from AI-related patents granted over the five years ending in the current year (detailed description in Section 3.6). Specifically, it refers to general AI exposure (AI^{EXP}) in columns (1) and (3) and AI net complementarity exposure (AI^{COMP}) in columns (2) and (4). All specifications incorporate the lagged occupational controls, including the average age, number of children, and educational attainment, and the share of female and married respondents, occupation fixed effects and year fixed effects. Standard errors are double clustered by occupation and year. Asterisks denote significance levels (***=1%, **=5%, *=10%).

Dep Var	ΔWeekly Work Hours$_{o,t}$			
	Female - Male		Young - Old	
	(1)	(2)	(3)	(4)
$AI^{EXP}_{o,t-1}$	0.039**		0.041*	
	(2.52)		(1.88)	
$AI^{COMP}_{o,t-1}$		0.064**		0.064*
		(2.19)		(2.00)
Occupation Controls	Yes	Yes	Yes	Yes
Occupation FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Observations	3,376	3,376	4,573	4,573
R^2	0.138	0.138	0.108	0.108
Adjusted R^2	0.057	0.057	0.041	0.041
Sample Mean of Dep Var	-4.55		-3.65	

Table 6: AI Technology Complementarity vs. Workday and Wage

The table reports weighted linear regression results that estimate the effect of occupational AI net complementarity on work hours and wages at the individual level. The ATUS survey sample and occupation classification are the same as in Table 3. The dependent variables are weekly work hours in columns (1)–(3) and the natural logarithm of hourly wages in 2023 dollars in columns (4)–(6). The main explanatory variable, AI^{COMP}, is AI net complementarity measure in percentile rank at the occupation-year level, calculated from AI-related patents granted over five years ending in the current year (detailed description in Section 3.6). All specifications control for individual characteristics, including age, the number of children, and a series of indicator variables for gender, educational attainment, marital status, and race, and fixed effects at the following levels: state × year, industry × year, year-month and day-of-week. Columns (2) and (5) additionally control for other occupational exposure measures, including robot exposure (Webb, 2019), routine task index (RTI) (Autor and Dorn, 2013), and offshorability exposure (Firpo et al., 2011; Autor and Dorn, 2013), all in percentile ranks. Columns (3) and (6) include occupation fixed effects, which subsume occupation-level controls. Standard errors are double clustered by occupation and state. Asterisks denote significance levels (***=1%, **=5%, *=10%).

Dep Var	Weekly Work Hours$_{i,o,t}$			100 × Log(Hourly Wage \$)$_{i,o,t}$		
	(1)	(2)	(3)	(4)	(5)	(6)
$AI^{COMP}_{o,t-1}$	0.056***	0.043***	0.026**	0.339***	0.244***	0.046**
	(4.10)	(3.77)	(2.22)	(7.49)	(4.74)	(2.13)
Individual characteristics	Yes	Yes	Yes	Yes	Yes	Yes
Other Occupation Exposure	No	Yes	No	No	Yes	No
Occupation FE	No	No	Yes	No	No	Yes
State × Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Industry × Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Year-Month FE	Yes	Yes	Yes	Yes	Yes	Yes
Day-of-Week FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	124,059	124,059	124,059	104,356	104,356	104,356
R^2	0.275	0.276	0.285	0.482	0.492	0.552
Adjusted R^2	0.262	0.264	0.270	0.471	0.481	0.542

Table 7: Exposure to AI Surveillance Technology and Workday

The table reports the weighted linear regression results based on individual responses of salaried employees working remotely in the ATUS survey from 2015 to 2023. The occupation classification is the same as in Table 3. Remote workers are defined as those in occupations with a work-from-home (WFH) feasibility index from Dingel and Neiman (2020) equals one. The dependent variable is weekly work hours. The main explanatory variable, AI_o^{SUR}, is AI surveillance exposure at the occupation level (detailed description in Section B.4 of Online Appendix) and transformed to percentile ranks (e.g., Autor and Dorn, 2013; Webb, 2019). Specifically, it refers to general AI surveillance exposure in column (1) and decomposed AI surveillance exposure in column (2)-(4) as specified in the third row. *POST* dummy equals one for the years since 2020. All specifications control for individual characteristics, including age, the number of children, and a series of indicator variables for gender, educational attainment, marital status, and race, and fixed effects at the following levels: occupation, state × year, industry × year, year-month, and day-of-week. Standard errors are double clustered by occupation and state. Asterisks denote significance levels (***=1%, **=5%, *=10%).

Sample	Weekly Work Hours$_{i,o,t}$			
Indep Var	Overall	Direction	Evaluation	Discipline
	(1)	(2)	(3)	(4)
$AI_o^{SUR} \times POST_t$	0.044**	0.043**	0.045**	0.032
	(2.15)	(2.03)	(2.25)	(1.48)
Individual characteristics	Yes	Yes	Yes	Yes
Occupational FE	Yes	Yes	Yes	Yes
State × Year FE	Yes	Yes	Yes	Yes
Industry × Year FE	Yes	Yes	Yes	Yes
Year-Month FE	Yes	Yes	Yes	Yes
Day-of-Week FE	Yes	Yes	Yes	Yes
Observations	10,238	10,238	10,238	10,238
R^2	0.509	0.509	0.509	0.509
Adjusted R^2	0.458	0.458	0.458	0.458

Table 8: AI Exposure and Firm Operating Performance

The table reports the linear regression results that examine the effects of occupational AI exposure on the operating performance of Compustat public firms from 2008 to 2023. The dependent variables are 100 times the return over assets (ROA) defined as operating income before depreciation over total assets in column (1)–(4) and labor productivity defined as sales over employment (in \$000) in column (5)–(8). The main explanatory variable is the annual workforce AI exposure at the firm level, calculated as employment-weighted averages of the corresponding occupation-level measures. Specifically, it refers to general AI exposure ($AI^{EXP}_{i,t-1}$) in columns (1) and (5), AI net complementarity exposure ($AI^{COMP}_{i,t-1}$) in columns (2) and (6), weekly work hours predicted by $AI^{EXP}_{i,t-1}$ in columns (3) and (7) and weekly work hours predicted by $AI^{COMP}_{i,t-1}$ in columns (4) and (8). Section 3.6 provides detailed descriptions of occupational AI exposure measures. Predicted work hours are based on estimates of individuals from column (3) of Table 3 and 6 and then aggregated to the occupation level using ATUS survey weights. Employment at the occupation-firm-year level is derived from LinkedIn data. All regressions control for firm attributes including sales in 2023 dollars (natural logarithm), Tobin's Q, market leverage, capital expenditure over the beginning-of-year assets, R&D expenditure over assets, net fixed assets over assets, firm fixed effects and year fixed effects. Standard errors are clustered at the firm level. Asterisks denote significance levels (***=1%, **=5%, *=10%).

Dep Var	$ROA_{i,t}\times100$				$Sales/Employment_{i,t}$			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$AI^{EXP}_{i,t-1}$	0.048***				0.009***			
	(3.59)				(5.67)			
$AI^{COMP}_{i,t-1}$		0.053***				0.008***		
		(3.05)				(4.54)		
$\widehat{Work\_Hour}^{EXP}_{i,t-1}$			0.258***				0.017***	
			(4.62)				(2.95)	
$\widehat{Work\_Hour}^{COMP}_{i,t-1}$				0.256***				0.017***
				(4.61)				(2.94)
Firm Attributes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	21,903	21,903	21,903	21,903	21,639	21,639	21,639	21,639
R^2	0.738	0.737	0.738	0.738	0.872	0.871	0.871	0.871
Adjusted R^2	0.701	0.701	0.702	0.702	0.854	0.854	0.854	0.854

Table 9: AI Exposure and Employee Satisfaction: Within Firm-Year

The table presents estimates from linear regressions examining the effects of occupational AI exposure on employee ratings at the occupation (o) × firm (i) × year (t) level. The occupation classification is the same as in Table 3. The sample covers data from private and public firms in the Glassdoor database ranging from 2008 to June 2023. The dependent variables are: 100 times the overall satisfaction rating in columns (1), the Work-Life Balance (WLB) ratings in columns (2), the average annual salary in 2023 dollars (in $000) obtained from Revelio in columns (3), and the number of complaints (negatively-toned reviews) mentioning surveillance in column (4) and mentioning employment risk in column (5). The main explanatory variable represents AI exposure measures at the occupation-year level, expressed in percentile ranks, and is based on AI-related patents granted in a five-year window ending in the current year (detailed description in Section 3.6). All specifications include lagged employee review counts, the average seniority and remote work index of the occupation-by-firm cohort from Revelio, and fixed effects at the following levels: occupation, firm×year. Standard errors are clustered by occupation. R^2 presents Pseudo R^2 for Poisson regressions. Asterisks denote significance levels (***=1%, **=5%, *=10%).

Dep Var	$\text{Rating}_{o,i,t}$			$\text{No.Complaints}_{o,i,t}$	
	Overall	WLB	$\text{Salary}_{o,i,t}$	Surveillance	Employment Risk
	(1)	(2)	(3)	(4)	(5)
$AI^{EXP}_{o,t-1}$	-0.052***	-0.049***	0.016**	0.001**	0.001
	(-2.94)	(-3.08)	(2.08)	(2.10)	(1.12)
Marginal Effect				0.020%	0.011%
Model	OLS	OLS	OLS	Poisson	Poisson
Cohort Controls	Yes	Yes	Yes	Yes	Yes
Occupation FE	Yes	Yes	Yes	Yes	Yes
Firm × Year FE	Yes	Yes	Yes	Yes	Yes
Observations	226,905	217,200	226,905	157,765	88,702
R^2	0.274	0.264	0.743	0.365	0.267
Adjusted R^2	0.216	0.203	0.723		

Table 10: AI Exposure and Workday: In Relation to Competition

The table reports the weighted linear regression results that estimate the heterogeneity effect of AI on work hours sorted by labor market and product market competition. The ATUS sample and occupation classifications follow Table 3. The dependent variable is weekly hours spent on market work. The main explanatory variable, AI^{COMP}, is AI net complementarity exposure in percentile rank at the occupation-year level, calculated from AI-related patents granted over five years ending in the current year (detailed description in Section 3.6). Two proxies for firms' labor market competition at the state-occupation level are specified: the employment concentration across firms measured by the Herfindahl-Hirschman Index (HHI) in columns (1) and firms' talent retention pressure (TRP) in columns (2). The state-occupation-level employment HHI is calculated using the employment and state information of public firms from LinkedIn. TRP is calculated as the job vacancy-to-employment ratio (V/E) at the state-occupation level, scaled by each occupation's cognitive skill scores, using job vacancy data from Burning Glass, employment data from the OEWS, and cognitive skill scores from O*NET. *I(Worker Power vs. Firm)* is the indicator of workers' bargaining power over firms that equals one if the measure for firms' labor market competition is in the top quartile (i.e., the employment HHI in the bottom quartile or TRP in the top quartile), and zero otherwise. Two proxies represent firms' product market power at the industry-level: product similarity in column (3) and product market concentration HHI in column (4), derived from firm-level scores from Hoberg and Phillips (2016) weighted by Compustat sales. *I(Firm Power vs. Consumer)* is an indicator of firms' product market power relative to consumers, which equals one if the product similarity is in the bottom quartile or the product HHI is in the top quartile, and zero otherwise. All specifications incorporate individual-level controls, including age, the number of children, and a series of indicator variables for gender, educational attainment, marital status, and race, and fixed effects at the following levels: occupation, state × year, industry × year, year-month and day-of-week. Standard errors are double clustered by occupation and state. Asterisks denote significance levels (***=1%, **=5%, *=10%).

Dep Var	Weekly Work Hours$_{i,o,t}$			
	Labor Market Competition		Product Market Competition	
Factor	HHI	TRP	Similarity	HHI
	(1)	(2)	(3)	(4)
$AI^{COMP}_{o,t-1}$	0.031**	0.080***	0.025**	0.028**
	(2.38)	(2.85)	(2.04)	(2.36)
× *I(Worker Power vs. Firm)*$_{o,s,t-1}$	-0.015*	-0.066**		
	(-1.86)	(-2.23)		
× *I(Firm Power vs. Consumer)*$_{j,t-1}$			-0.012	-0.031***
			(-1.50)	(-3.00)
Individual Characteristics	Yes	Yes	Yes	Yes
Occupational FE	Yes	Yes	Yes	Yes
State × Year FE	Yes	Yes	Yes	Yes
Industry × Year FE	Yes	Yes	Yes	Yes
Year × Month FE	Yes	Yes	Yes	Yes
Day-of-Week FE	Yes	Yes	Yes	Yes
Observations	102,434	48,349	114,661	114,661
R^2	0.289	0.290	0.283	0.283
Adjusted R^2	0.272	0.272	0.268	0.268

Table 11: AI Exposure, New Hires and Separations

The table reports the weighted linear regression results that examine the effect of AI on hiring and separation at the occupation-state-year level from 2008-2023. The hiring and separation data is sourced from LinkedIn profiles of workers at public firms provided by Revelio. The occupation classification is the same as in Table 3. The dependent variable is 100 times the natural logarithm of new hires, all separations, voluntary separations, and involuntary separations in columns (1)–(4) and in columns (5)–(8). The main explanatory variable is AI exposure in percentile rank at the occupation-year level, calculated from AI-related patents granted over the five years ending in the current year (detailed description in Section 3.6). All models incorporate lagged controls at the occupation-by-state cohort level, including LinkedIn employment (natural logarithm), average seniority and remote work index from Revelio, and fixed effects and the following level: occupation and state × year. The regression is weighted by the underlying LinkedIn employment at the occupation-state-year level. Standard errors are double clustered by occupation and state. Asterisks denote significance levels (***=1%, **=5%, *=10%).

Dep Var	$100\times\text{Ln}(\text{Outcome}_{o,s,t})$							
	New Hires	Separations			New Hires	Separations		
		All	Voluntary	Involuntary		All	Voluntary	Involuntary
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$AI^{EXP}_{o,t-1}$	0.027	0.027	0.073**	0.009				
	(1.01)	(0.97)	(2.19)	(0.29)				
$AI^{COMP}_{o,t-1}$					0.061	0.061	0.144***	0.034
					(1.41)	(1.57)	(3.13)	(0.71)
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
State × Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	145,608	143,226	122,518	141,600	145,608	143,226	122,518	141,600
R^2	0.978	0.984	0.970	0.981	0.978	0.984	0.970	0.981
Adjusted R^2	0.978	0.984	0.970	0.981	0.978	0.984	0.970	0.981

Online Appendix

A. Optimal worker effort in a Principal-Agent Model

This model is a simple adaptation of Holmstrom and Milgrom (1987), aiming at illustrate the relation between agent "effort" (which maps to length of work day) and a set of factors including marginal productivity, effort observability, and bargaining power.

The output X_t follows a continuous-time stochastic process, affected by the agent's effort a and a noise term that is outside the control of the agent:

$$dX_t = \gamma a_t dt + \sigma dW_t, \tag{7}$$

where:

a_t is the agent's effort level ("working time" in our empirical setting) at time t, which is not directly observed by the principal. γ is the productivity parameter. σ represents the level of uncertainty in the noise term, and W_t is the standard Wiener process.

The principal is risk neutral with the following utility function V, which is the difference between the expected output γa_t and the compensation to the agent, $C_t = f(X_t)$:

$$V = \int_0^1 (\gamma a_t - f(X_t))dt \tag{8}$$

Effort, a_t, is not contractible and hence the compensation function relies on output which is a noisy representation of agent effort.

The agent is risk-averse with CARA utility with a risk-aversion coefficient of r, with a utility function depending on income C and leisure, and with a researvation utility of U_0. Assume the agent has one unit of time to allocate between work and leisure, his expected utility is $E[U(C, 1-a)]$. If we rule out the income effect of leisure for now, we assume that the U take the simple form of

$$U = \int_0^1 (C_t - \frac{1}{2} r Var(C_t) - \frac{1}{2} k a^2)dt \tag{9}$$

The principal solves the following optimization problem:

$$\begin{aligned} \text{Max}_{f(X_t)} \;\; V &= \int_0^1 (\gamma a_t - f(X_t))dt \\ s.t. E(U[f(X_t), a_t^*]) &\geq \underline{U} \; (Participation\ constraint) \\ a^* = \text{Argmax}_{a_t} \;\; E(U[f(X_t), a_t]) &\; (Incentive\ compatibility) \end{aligned} \tag{10}$$

Holmstrom and Milgrom (1987) shows that the optimal dynamic contract converges in the aggregate to a linear contract in the form of

$$C_t = \alpha + \beta X_t, \tag{11}$$

where β could be characterized as

$$\beta = \frac{1}{1 + kr\sigma^2} \tag{12}$$

Finally, the agent's effort level in response to the incentive is

$$a_t = \frac{\gamma}{k(1 + kr\sigma^2)} \tag{13}$$

In summary, equilibrium effort input is positively related to γ, the marginal productivity of effort; and negatively associated with k, the marginal cost of effort; r, the agent's risk aversion; and σ, the volatility of the noise in performance attribution to agent effort. Such comparative statics are robust with more general functional forms, though there is no closed-form solution.

In this simple model when agent's utility function is separable in consumption and leisure (see equation 9), a change in the agent's reservation utility (which is determined by her next best alternatives) does not affect the incentives and effort input. This will change with the relaxation of agent's utility function to a more general form, such as the constant elasticity of substitution (CES) utility function:

$$U(C, 1-a) = [\eta C^{\rho} + (1-\eta)(1-a)^{\rho}]^{\frac{1}{\rho}} - \frac{1}{2} r Var(C_t), \tag{14}$$

where $\eta \in (0, 1)$ is the relative preference for consumption and leisure, and $\rho < 1$ is the substitution parameter, or $\frac{1}{1-\rho}$, the elasticity of substitution between C and $1 - a$, is strictly positive.

Under this setup, the relation between a^* and $\underline{U}$ is not monotone. However, under reasonable parameters (e.g., agents are reasonably risk averse, and measuring performance is reasonably noisy), increasing $\underline{U}$ (because the agent has better outside opportunities due to bargaining power over their employer and the job market) tends to decrease work time. In addition, the following two conditions would each on its own serve as a sufficient condition for effort (work time) to shrink when $\underline{U}$ rises:

1. $\rho < 0$, i.e., consumption and leisure are strict complements.
2. $\underline{U}$ is sufficiently large, such that there is a limit on increasing β to agent the required utility due to agent's risk aversion.

Overall, because leisure is a normal good, the agent values leisure more when the agent's welfare improves. This force induces the agent to allocate more time into leisure from work, other things equal. The effect is stronger when agent risk aversion is high; performance measurement is noisy, complimentarity between consumption and leisure is high, and the agent has good alternatives (hence demands high reservation utility).

B. Documentation

B.1. Historical panel of O*NET data

The O*NET Data Collection Program currently makes updates to the O*NET Database quarterly, with a primary update occurring in the 3rd quarter (August) of each year. Prior to year 2015, the data was primarily updated once per year. To create a consistent annual panel of job tasks, we use the O*NET databases released each August from 2015 onward. For years prior to 2015, we select the data release closest to August, prioritizing those published between June and August when multiple versions are available in the same year. Table OA.1 of Online Appendix lists the O*NET data release we use to construct the annual panel of occupations' job tasks from 2000–2023.

B.2. Measure AI exposure at the occupation level using GPT

This section provides details on how we use OpenAI's GPT model to quantify AI exposure measures for occupations. Our practice was conducted on November 22, 2024 using the "gpt-4o-mini-2024-07-18" model with the "temperature" parameter set to 0.[30]

B.2.1. Prompt setup

GPT (Generative Pre-trained Transformer) architecture, developed by OpenAI, uses a transformer design with self-attention mechanisms for advanced contextual understanding. Pre-trained on vast datasets, it is highly proficient in processing and analyzing text.

We use GPT to classify the impact of AI patents on occupations due to its ability to identify complex relationships and nuances in language. Specifically, we define a prompt, which serves as a clear instruction or context-setting input that shapes the model's output, as following and apply it to a given patent-occupation combination in our sample:

You are a labor economist. Evaluate the extent to which a new AI patent substitutes or complements job tasks of a given occupation, and its impact on task completion time. Respond strictly in JSON format:

"overlap": [similarity_score], # Similarity between patent and tasks (1-10)

"label": [effect_label], # indicator of the impact of patent on tasks (-1 = substitute, 1 = complement, 0 = unrelated)

Include no text other than the JSON object.

In this prompt, we ask GPT to assume the role of a labor economist to classify the impact of a patent filing on a given occupation. The terms Patent Title and Patent Abstract are substituted by the title and abstract of a particular patent during the query. Similarly, Occupation Title and Tasks are substituted by the title and the combined text of all task statements of a particular occupation.

B.2.2. Example

We provide two examples of how GPT scores the overlap between an occupation and AI patent and identifies their substitute/complementarity relation.

Example 1)

Occupation: Retail Salespersons (SOC Code: 41-2031.00)

Task Statements: "Greet customers and ascertain what each customer wants or needs.| Open and close cash registers, performing tasks such as counting money, separating charge slips, coupons, and vouchers, balancing cash drawers, and making deposits.| Maintain knowledge of current sales and promotions, policies regarding payment and exchanges, and security practices.| Compute sales prices, total purchases and receive and process cash or credit payment.| Maintain records related to sales.| Watch for and recognize security risks and thefts, and know how to prevent or handle these situations.| Recommend, select, and help locate or obtain merchandise based on customer needs and desires.| Answer questions regarding the

[30] Temperature is a parameter in GPT model that controls the randomness and creativity of its responses. Setting the temperature to 0 makes the model consistently choose the most probable word.

store and its merchandise.| Describe merchandise and explain use, operation, and care of merchandise to customers.| Ticket, arrange and display merchandise to promote sales.| Prepare sales slips or sales contracts.| Place special orders or call other stores to find desired items.| Demonstrate use or operation of merchandise.| Clean shelves, counters, and tables.| Exchange merchandise for customers and accept returns.| Bag or package purchases, and wrap gifts.| Help customers try on or fit merchandise.| Inventory stock and requisition new stock.| Prepare merchandise for purchase or rental.| Sell or arrange for delivery, insurance, financing, or service contracts for merchandise.| Estimate and quote trade-in allowances.| Estimate cost of repair or alteration of merchandise.| Estimate quantity and cost of merchandise required, such as paint or floor covering.| Rent merchandise to customers."

Patent #1 Title: Use of product viewing histories of users to identify related products (Patent ID: 6912505, Year: 2005)

Patent Abstract: "Various methods are disclosed for monitoring user browsing activities that indicate user interests in particular products or other items, and for using such information to identify items that are related to one another. In one embodiment, relationships between products within an online catalog are determined by identifying products that are frequently viewed by users within the same browsing session (e.g., products A and B are related because a significant portion of those who viewed A also viewed B). The resulting item relatedness data is preferably stored in a table that maps items to sets of related items. The table may be used to provide personalized product recommendations to users, and/or to supplement product detail pages with lists of related products. In one embodiment, the table is used to provide session-specific product recommendations to users that are based on the products viewed by the user during the current browsing session."
GPT Overlap Score: 7. GPT Label: Complement.

Patent #2 Title: Payment transaction authentication system and method (Patent ID: 10755281, Year: 2020)

Patent Abstract: "This disclosure describes, in part, techniques for validating a payment transaction between a customer and a merchant via challenge questions. For instance, the method includes determining, by a payment processing system, a level of risk associated with a current payment transaction between the merchant and the customer; in response the level of risk being higher than a threshold, obtaining a query for the customer, wherein the query is based at least on the current payment transaction or one or more past transactions involving the customer; receiving, from a customer device associated with the customer, a response to the query; and validating the current payment transaction based on the response." **GPT Overlap Score: 6. GPT Label: Substitute.**

Example 2)
Occupation: "Office Clerks, General" (SOC Code: 43-9061.00)

Task Statements: "Operate office machines, such as photocopiers and scanners, facsimile machines, voice mail systems, and personal computers.| Answer telephones, direct calls, and take messages.| Communicate with customers, employees, and other individuals to answer questions, disseminate or explain information, take orders, and address complaints.| Maintain and update filing, inventory, mailing, and database systems, either manually or using a computer.| Compile, copy, sort, and file records of office activities, business transactions,

and other activities.| Review files, records, and other documents to obtain information to respond to requests.| Open, sort, and route incoming mail, answer correspondence, and prepare outgoing mail.| Compute, record, and proofread data and other information, such as records or reports.| Complete work schedules, manage calendars, and arrange appointments.| Type, format, proofread, and edit correspondence and other documents, from notes or dictating machines, using computers or typewriters.| Inventory and order materials, supplies, and services.| Deliver messages and run errands.| Collect, count, and disburse money, do basic bookkeeping, and complete banking transactions.| Complete and mail bills, contracts, policies, invoices, or checks.| Process and prepare documents, such as business or government forms and expense reports.| Monitor and direct the work of lower-level clerks.| Prepare meeting agendas, attend meetings, and record and transcribe minutes.| Train other staff members to perform work activities, such as using computer applications.| Count, weigh, measure, or organize materials.| Make travel arrangements for office personnel.| Troubleshoot problems involving office equipment, such as computer hardware and software."

Patent #1 Title: Generative summaries for search results (Patent ID: 11769017, Year: 2023)

Patent Abstract: "At least selectively utilizing a large language model (LLM) in generating a natural language (NL) based summary to be rendered in response to a query. In some implementations, in generating the NL based summary additional content is processed using the LLM. The additional content is in addition to query content of the query itself and, in generating the NL based summary, can be processed using the LLM and along with the query content or even independent of the query content. Processing the additional content can, for example, mitigate occurrences of the NL based summary including inaccuracies and/or can mitigate occurrences of the NL based summary being over-specified and/or under-specified."
GPT Overlap Score: 7. GPT Label: Complement.

Patent #2 Title: Method and system for synchronizing databases automatically and periodically (Patent ID: 10936623, Year: 2021)

Patent Abstract: "Through a first processing thread, a first database is accessed via a first API to retrieve a list of event objects of the first database. Through a second processing thread, for each of the event objects, participant identifiers (IDs) are determined from the event object. For each of the participant IDs, a domain ID is extracted from the participant ID. A list of one or more entity objects are identified based on the domain ID, where the entity objects are stored in a second database such as a task database storing and managing many tasks. At least one attribute of at least one of the entity objects is modified based of the participant ID and the domain ID, which generates a modified entity object. Through a third processing thread, any event objects that have been modified are transmitted to the second database via a second API over the network."
GPT Overlap Score: 8. GPT Label: Substitute.

B.2.3. Validating GPT-generated AI general exposure

Generative Large Language Models, such as GPT, provide improved textual analysis approaches over non-generative methods, mainly because that they enable expressing a task

through natural language and exhibit more sophisticated reasoning abilities (de Kok, 2025). However, the black-box nature of these models poses challenges to the validation of the measures created by them. Here, we apply a non-generative natural language processing method to calculate a comparative variable to our overlap variable generated by GPT. Specifically, following Kogan et al. (2023), we employ a combination of word embedding and term-frequency-inverse-document-frequency (TF-IDF) approach to calculate the similarity between the text description of an occupation and the abstract of a patent. Then, we aggregate the similarity score at the occupation-year level to represent the time-varying relevance of AI to each occupation's tasks. Finally, we compare the TF-IDF ee similarity score to the GPT-generated AI exposure score.

The specific procedure is as follows. First, we pre-process each text portion of the task description of each SOC 8-digit occupation and patent abstracts by removing non-alphabetic characters, lowercasing all text, removing all stopwords listed in the sources in Kogan et al. (2023), and retaining lemmatized versions of nouns and verbs only. Next, we represent each word of a text as a 100-dimensional vector using the word vectors provided by Pennington et al. (2014). The word vectors are numerical representations of word meanings that can effectively capture pairwise distances between words based on co-occurrence probabilities (Kogan et al., 2023). Then, to measure the document similarity between an occupation task description and a patent abstract, we construct a document-level vector, which is a weighted average of the set of word vectors in each task description or patent abstract text. We use TF-IDF to weigh each word vector, which gives higher weights for terms that occur more frequently in a document and lower weights for terms that occur commonly across many documents (Kogan et al., 2023). Finally, we calculate the cosine similarity between the task description of each occupation and a patent abstract, each represented as a document vector, to measure the relevance of the AI patent to the tasks performed by the occupation.

We aggregate the TF-IDF similarity scores from the SOC 8-digit occupation by patent level to the *occ*1990*dd* occupation by year level following the procedures outlined in Section 3.6.2. The TF-IDF score and the GPT-based AI exposure score ($AI_{o,t}^{EXP}$) constructed using the same patents show a high correlation of 0.83, demonstrating a strong alignment between the two measures and confirming the robustness of the GPT-derived approach.

B.2.4. Validating GPT-generated AI net complementarity exposure

To study the wage effects of AI, Kogan et al. (2023) use ChatGPT4 released in March 2023 to identify whether AI is a substitute or complement to occupation tasks using a question-based approach. Specifically, they ask ChatGPT whether AI's is able to perform specific job tasks with or without human intervention. This approach yields time-invariant measures of the occupation's exposure to AI substitution and AI complementarity.

Kogan et al. (2023) do not report the AI exposures at the occupation level. However, Table A8–9 in their Online Appendix provide different components of AI exposure-related earnings changes for occupations with the highest AI substitution (or complementarity) exposure at the SOC 6-digit occupation level. We validate our AI Net complementarity exposure ($AI_{o,t}^{COMP}$) by comparing it to those wage growth components. To summarize, $AI_{o,t}^{COMP}$ in 2023 exhibits a strong negative correlation of -0.59 with the wage growth attributed to the substitution effect of AI (column (3) in Table A8–9) documented in Kogan et al. (2023)). In contrast, $AI_{o,t}^{COMP}$ shows a positive correlation of 0.47 with wage growth related to labor-complementing effects

(column (4) of in Table A8–9) and a correlation of 0.60 with the overall wage growth of AI (column (6)). This underscores a strong consistency between the methods, validating the reliability of our AI net complementarity measure.

B.3. Measure generative AI exposure at the occupation level using GPT

This section provides details on how we use a GPT model to replicate the exposure to Generative AI of each task following Eisfeldt et al. (2023) and to distinguish the impact of substitution and complementarity of Generative AI on each task following Kogan et al. (2023). The categorization was conducted on November 12, 2024 using the "gpt-4o" model with the GPT "temperature" parameter set to 0. The job task descriptions of occupations are obtained from the O*NET 27.0 database released on August 1, 2022. Using the task statement, we generated two output variables for each of the 19,267 tasks including (i) to which extent the task is exposed to Generative AI technologies or not (ii) whether it is substituted or complemented by Generative AI technologies.

The replicated generative AI exposure was used for the initial analyses; however, following the release of the original measure by Eisfeldt et al. (2023), we replaced it with the original to ensure consistency and comparability.

B.3.1. Prompt setup

We define a prompt as following and apply it to each job task in our sample:

"Generate two outcomes in the exact format of '[val1, val2, val3], [label]' based on the following instructions."

"First Task: Pretend you are a labor economist evaluating the extent to which Generative AI (specifically ChatGPT) might substitute or complement a job task of an occupation. "

"The output must be exactly a list of numbers in this format: [val1, val2, val3], where: - val1 is ChatGPT's substitute score (1-10)," "- val2 is its complement score (1-10)," "- val3 is a label (-1 = substitute, 1 = complement, 0 = unrelated) indicating if ChatGPT primarily complements or substitutes the job task."

"Second Task: For the second task, use the following Context for Evaluation and Exposure Rubric to label a given occupation task with one of the labels (E0, E1, E2, or E3) based on its exposure to LLM capabilities."

"The output must be exactly in this format: [label] that best describes the task's exposure to the LLM."

"Context for Evaluation: Assume access to the most powerful OpenAI large language model (LLM). This model can complete tasks involving text input and output, as long as the context can be captured in 2000 words. However, it cannot retrieve up-to-date facts from the past year unless provided in the input. Assume you are a worker with average expertise, using the LLM along with other software or hardware tools specified in the task. You also have commonly available technical tools (e.g., microphone, speakers) but no other physical materials. Your goal is to label tasks according to the rubric below, ensuring equivalent quality (i.e., a reviewer cannot distinguish whether a human completed it independently or with LLM assistance). If you are unsure how to judge time savings, consider if the described tools cover the majority of the subtasks."

"Exposure Rubric:"

"- E1 - Direct Exposure: Label tasks as E1 if direct access to the LLM (e.g., via ChatGPT or OpenAI playground) alone can reduce task time by at least half while maintaining quality. Examples include:" " - Writing and transforming text/code," " - Editing text/code as specified,"

" - Writing code for tasks previously done manually," " - Translating text," " - Summarizing medium-length documents," " - Providing document feedback," " - Answering questions about a document," " - Generating or answering questions," " - Writing or responding to emails (including negotiation if via text)," " - Maintaining written records," " - Preparing general training materials, and" " - Informing others through written or spoken formats."

"- E2 - Exposure by LLM-powered Applications: Label tasks as E2 if the LLM alone may not halve the time required, but additional software built on the LLM could. Examples include:" " - Summarizing documents longer than 2000 words and answering questions on them," " - Retrieving recent/specialized information from the internet or organization data," " - Making recommendations based on data," " - Analyzing written information for decisions," " - Preparing specialized training materials, and" " - Maintaining complex databases."

"- E3 - Exposure with Image Capabilities: Label tasks as E3 if the combination of the LLM and an image-processing system (capable of viewing, captioning, and creating images, but not video) significantly reduces task time. Examples include:" " - Reading text from PDFs," " - Scanning images," " - Creating or editing digital images based on instructions (realistic but not highly detailed)."

"- E0 - No Exposure: Label tasks as E0 if none of the above criteria apply, and no clear reduction in task time by half is achieved. Examples include:" " - Tasks requiring significant human interaction (e.g., in-person demonstrations)," " - Tasks requiring precise physical measurements or detailed visual review," " - Decisions impacting human livelihood (e.g., hiring, grading)," " - Tasks legally requiring a human," " - Tasks already completed efficiently with existing (non-LLM) technology, and" " - When in doubt, default to E0."

B.3.2. Variable construction

Task scoring By applying the prompt, we categorize the Generative AI exposure, $GenAI_j$ of a given task j into one of the following three categories based on the GPT output in the second task of the prompt:

- Direct Exposure ($GenAI_j = 1$): if ChatGPT enables a task to be completed in less than half the usual time, maintaining the same quality.
- Plus-Overlay Exposure ($GenAI_j = 0.5$): if ChatGPT alone cannot cut task time by half, but the addition of complementary software leveraging its functionality could achieve this efficiency without sacrificing quality.
- No Exposure ($GenAI_j = 0$): if ChatGPT neither reduces task time by half with comparable quality nor produces results of adequate quality.

Meanwhile, we classify a given task j as being substituted or complemented by Generative AI into one of the following three classifications based on the GPT output "label" in the first task of the prompt:

- Substitute ($GenAI_j^{COMP} = -1$): if ChatGPT primarily substitutes a job task.
- Complement ($GenAI_j^{COMP} = 1$) if ChatGPT primarily complements a job task.
- Unrelated ($GenAI_j^{COMP} = 0$) if ChatGPT is irrelevant to a job task.

Aggregation to the Occupation-Level We next aggregate tasks' exposures to Generative AI to the SOC 8-digit occupation level. Following Eisfeldt et al. (2023), we calculate the Generative AI exposure ($GenAI_o$) of a given occupation as the share of the total number of tasks for each occupation that have either a direct or "plus-overlay" exposure to Generative AI. We calculate Generative AI - Net complementarity ($GenAI_o^{COMP}$) for each SOC 8-digit occupation by taking the equal-weighted average of $GenAI_j^{COMP}$ across all tasks associated with that occupation. Next, we aggregate SOC 8-digit occupation codes to $occ1990dd$ codes following the procedures outlined in Section 3.6.2.

B.3.3. Validation

We validate $GenAI_o^{COMP}$ by comparing it to Table A8–9 of Kogan et al. (2023). Table A8–9 of Kogan et al. (2023) reports the predicted wage growth attributed to different components of AI exposure of occupations with the highest exposure to labor-complementing and labor-substituting potential of AI at the SOC 6-digit level. We find that $GenAI_o^{COMP}$ has a correlation of -0.42 with the wage growth attributed to the labor-substituting potential of AI (column (3) of Table A8–9 in Kogan et al. (2023)), and a correlation of 0.31 with that of labor-complementing (column (4)) and a correlation of 0.44 with the total wage growth of AI (column (6)), documented by Kogan et al. (2023)).

B.4. Measure AI surveillance exposure at the occupation level using GPT

This section provides details on how we use a GPT model to quantify the exposure to AI surveillance of each task. We conducted this categorization on March 28, 2025 using the "gpt-4o-2024-11-20" model with the GPT "temperature" parameter set to 0. The job task descriptions of occupations are obtained from the O*NET 28.0 database released on August 1, 2023. Using the task statement, we generated four output variables for each of the 19,280 tasks including to which extent the task is exposed to overall and decomposed AI surveillance technologies.

B.4.1. Prompt setup

We define a prompt as following and apply it to each job task in our sample:

*"As a labor economist, assess AI's ability to improve **control efficiency** by better tracking and evaluating workers' performance, effort, and compliance based on three perspectives of Algorithmic Control: Direction, Evaluation, and Discipline."*

*" **Context for Assessment:**"*

*"1. **Algorithmic Direction** – AI guides or restricts workers' actions to align with goals."*

*" - **Recommending:** Prompts workers to align decisions with predefined goals."*

*" - *Example:* AI recommends optimal scheduling based on data analysis."*

*" - **Restricting:** Limits access to information or constrains behavior."*

*" - *Example:* AI restricts information or modifies behavior in online communities."*

*""2. **Algorithmic Evaluation** – AI monitors and assesses performance through data analysis."*

*" - **Recording:** Tracks behaviors and provides real-time feedback."*

*" - *Example:* AI logs work speed and accuracy for reviews."*

*" - **Rating:** Aggregates data (e.g., ratings, rankings) to evaluate productivity and predict performance."*

*" - *Example:* AI ranks employees based on task completion rates."*

*""3. **Algorithmic Discipline** – AI enforces compliance and incentivizes workers via automation and rewards."*

*" - **Replacing:** Automatically removes or reassigns underperforming workers."*

*" - *Example:* AI flags low-rated workers for reassignment."*

*" - **Rewarding:** Provides dynamic rewards or gamifies tasks to increase engagement."*

*" - *Example:* AI gives real-time rewards for task completion."*

*" **Output Format:**"*

"One exact response per job task in the format: '[v1,v2,v3],[label]'."

*"- **[v1, v2, v3]**: Scores from 1–10 for direction (v1), evaluation (v2), and discipline (v3), where 10 indicates maximum potential for AI to improve control.*

*"- **label**: Composite AI control score (1–10) reflecting all three aspects."*

B.4.2. Variable construction

Task scoring By applying the prompt, the GPT model returns responses three scores– $v1$, $v2$, and $v3$– range from 1 (low potential) to 10 (high potential) for each dimension of AI

surveillance including direction, evaluation, and discipline, along with a composite score (*label*) representing a general AI surveillance score. We drop tasks that GPT could not classify, leading to 19,273 tasks.

Given upward biases in raw scores (e.g., means: $v1 = 6.9$, $v2 = 8.0$), we convert these continuous scores into binary exposure indicators using their respective sample means as thresholds. Specifically, we define:

- General AI surveillance exposure: $AI_j^{SUR} = 1$ if $label > 7$;
- Direction component: $AI_j^{SUR} - Direction = 1$ if $v1 > 7$;
- Evaluation component: $AI_j^{SUR} - Evaluation = 1$ if $v2 > 8$;
- Discipline component: $AI_j^{SUR} - Discipline = 1$ if $v3 > 6$.

Aggregation to the occupation-level We then aggregate these binary task-level AI surveillance indicators to the SOC 8-digit occupation level. Following Eisfeldt et al. (2023), occupation-level exposures (AI_o^{SUR} and its decomposed measures) are computed as equal-weighted averages of corresponding task-level indicators across all tasks associated with each SOC 8-digit occupation. Finally, SOC 8-digit occupation codes are aggregated to $occ1990dd$ using the procedure outlined in Section 3.6.2.

B.4.3. Example

Table OA.8 of the Online Appendix lists top occupations grouped by general AI surveillance exposure. On the top of the list are travel agent, dispatchers, bookkeeping clerks, medical records specialists, etc., while occupations with the lowest AI surveillance exposure include clergy, dentists, and artists.

B.5. Measure employee sentiment toward surveillance and employment risk

This section outlines how we quantify firm-occupation-year-level complaint measures about surveillance and employment risk using Glassdoor employee reviews.

B.5.1. Measurement

Each Glassdoor review includes two text fields: "Pros" and "Cons." We concatenate these fields for each English-language review to compute the net sentiment score of all texts of a review. Then, we merge frequently occurring bigrams and trigrams, such as "work-life balance."

To identify text related to each of the two topics (i.e., surveillance and employment risk), we construct keyword lists for both topics. Using Gensim's Word2Vec, we train a word embedding model on the entire corpus of employee reviews. For surveillance, we use "performance," "monitor," "surveillance," and "quota." For employment risk, we seed the model with initial keywords: "layoff," "unemployment," "downsizing," and "termination." We expand each list by retrieving words with cosine similarity greater than 0.5 to the seed terms. After manual inspection, the final surveillance keyword list includes 71 keywords (e.g., "metric," "evaluations") and the employment risk keyword list contains 66 keywords (e.g., "displacement," "recession").

To calculate sentiment toward each topic for each employee review, we first extract two sets of sentences: those containing any surveillance keyword and those containing any employment risk keyword. We then apply a transformer-based sentiment analysis model to each set of sentences separately to assess topic-specific sentiment. Specifically, we use the "cardiffnlp/twitter-roberta-base-sentiment" model, which outputs "positive," "neutral," and "negative" scores that sum to 1 for a given text. A review's sentiment score for each topic is computed as the positive score minus the neutral score minus the negative score. For each topic, we classify reviews with sentiment scores below the 20th percentile as "Negative." Finally, for each firm-occupation-year, we compute the two complaint measures: the number of reviews classified as "Negative" toward surveillance and the number of reviews classified as "Negative" toward employment risk.

B.5.2. Example

Table OA.10 the Online Appendix lists examples of complaints about surveillance and employment risk.

B.6. Map leisure activities to personal utility parameters

This section details how we mapped ATUS leisure activities to the 126 personal utility parameters of life aspects identified by Benjamin et al. (2025). This mapping was performed on July 23, 2025, using the "claude-opus-4-20250514" model developed by Anthropic, with the temperature parameter set to 0. For each life aspect, we generated eight output variables: (i) an overall score of leisure's impact, (ii) a one-sentence rationale for the score in (i), and (iii) subcategory-specific impact scores for leisure domains such as recreation and eating, sleeping, and personal care (ESP).

B.6.1. Prompt setup

We use the following prompt template to instruct the Claude model to evaluate the impact of increased leisure hours on each life aspect:

"You are a reseracher in labor economics who strictly follows instructions and provides only valid JSON answers."

*"Evaluate the direct impact of an increase in a worker's **leisure hours (screen time, social time, recreation, eating/sleeping/personal care, and hobbies)** at the expense of **work hours** on the well-being aspect of **<Aspect>**, holding hours for other activities (i.e., education, civic activities, medical care, job search, home production, child care) constant. Do not assume relationships without justification. Also consider the consequences of both increased leisure and decreased work hours. Focus strictly on the direct impact of leisure hours on the aspect, not indirect impacts. Use 0 labels/scores when appropriate."*

*"Provide your answer strictly in the following JSON format. Respond **only** with valid JSON and **no additional text**. Make sure every key-value pair ends with a comma, except the last one. Respond only with syntactically valid JSON that can be parsed by 'json.loads()'."*

*"**leisure_score**: [float], // A number from -1 to 1 representing the direct effect of increased leisure hours at the expense of work hours on **<Aspect>**. Use 0 if there is no known or meaningful relationship.*

*"**leisure_reason**: [str], // A one sentence reason for leisure_score.*

*"**leisure_screen_score**: [float], // A number from -1 to 1 representing the direct effect of increased screen-based leisure (e.g., gaming, TV) at the expense of work hours on **<Aspect>**. Use 0 if there is no known or meaningful relationship.*

*"**leisure_nonscreen_score**: [float], // A number from -1 to 1 representing the direct effect of increased non-screen leisure hours (e.g., recreation, socializing, dining well, sleep) at the expense of work hours on "<Aspect>". Use 0 if there is no known or meaningful relationship.*

*"**leisure_recreation_score**: [float], // A number from -1 to 1 representing the direct effect of increased recreational leisure hours (e.g., relaxing, music, traveling) at the expense of work hours on "<Aspect>". Use 0 if there is no known or meaningful relationship.*

*"**leisure_social_score**: [float], // A number from -1 to 1 representing the direct effect of increased socializing time at the expense of work hours on "<Aspect>". Use 0 if there is no known or meaningful relationship.*

*"**leisure_ESP_score**: [float], // A number from -1 to 1 representing the direct effect of increased hours for eating, sleeping, and personal care at the expense of work hours on "<Aspect>". Use 0 if there is no known or meaningful relationship.*

*"**leisure_other_score": [float], // A number from -1 to 1 representing the direct effect of increased other leisure hours (hobbies, reading, sports) at the expense of work hours on "<Aspect>". Use 0 if there is no known or meaningful relationship.*

B.6.2. Variable construction

Applying the prompt yields scores ranging from -1 (negative impact) to 1 (positive impact), representing the estimated extent to which each of the 126 life aspects is impacted by leisure. We then compute the utility parameter for each leisure category as a weighted sum of the relative marginal utilities of the life aspects (as reported by Benjamin et al. (2025)), where the weights are the LLM-generated impact scores reflecting each leisure category's impact on those aspects.

Due to the size of the resulting table, we provide a link to the table that contains 126 life aspects, LLM-generated prompt outputs of eight variables, and the computed utility parameters for each leisure category: https://www.dropbox.com/scl/fi/muoey2ltc2vmo8ion7jkd/claude_aspects_overleaf.xlsx?rlkey=4arh4nisie8z4ntfcn9b6fxlp&st=cnrn65q2&dl=0.

C. Figures

Figure OA.1. AI Patents Over Time

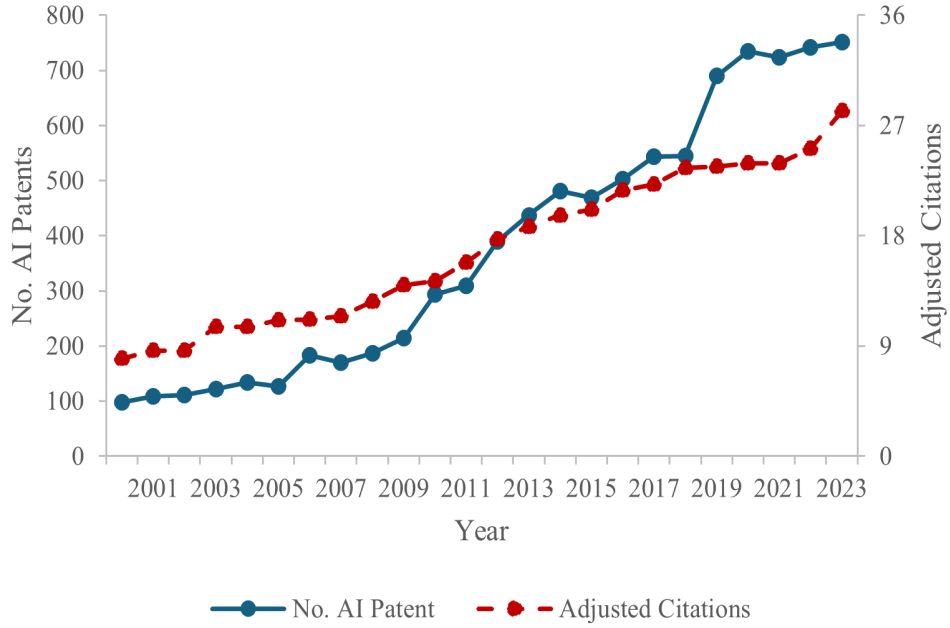

The figure summarizes the top 1% high-impact AI patents every year from 2000 - 2023 based on the adjusted forward citations that are used to construct AI exposure measures. Following Kogan et al. (2017), a patent's adjusted forward citations are calculated as its raw citation count divided by the average citation count of AI patents granted in the same year-quarter and CPC subclass. The blue line represents the total number of AI patents granted each year, while the red line depicts the average adjusted forward citations for that year's cohort.

Figure OA.2. AI Exposure vs. General Technology Exposure

A: Felten et al. AI Exposure

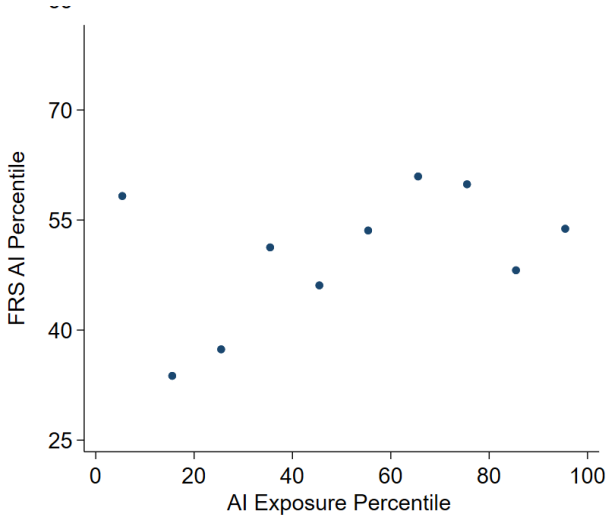

A: Webb AI Exposure

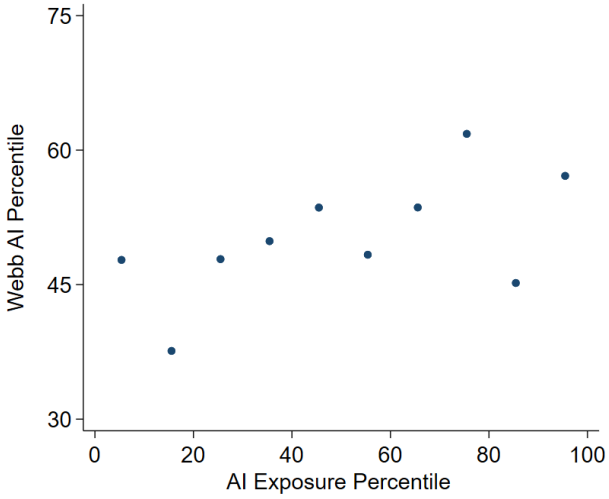

C: Webb Robot Exposure

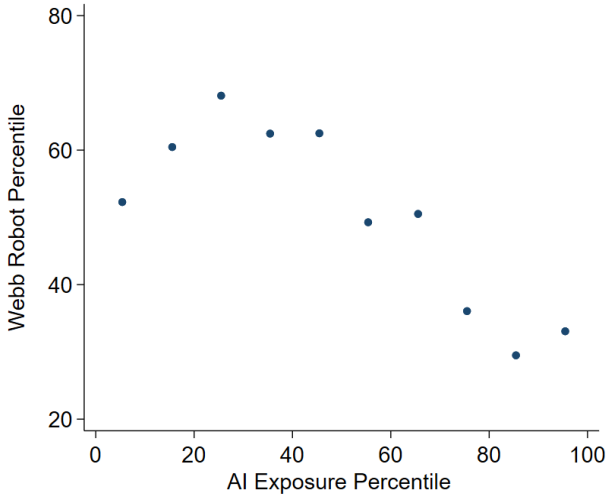

D: Routine-Task Intensity (RTI)

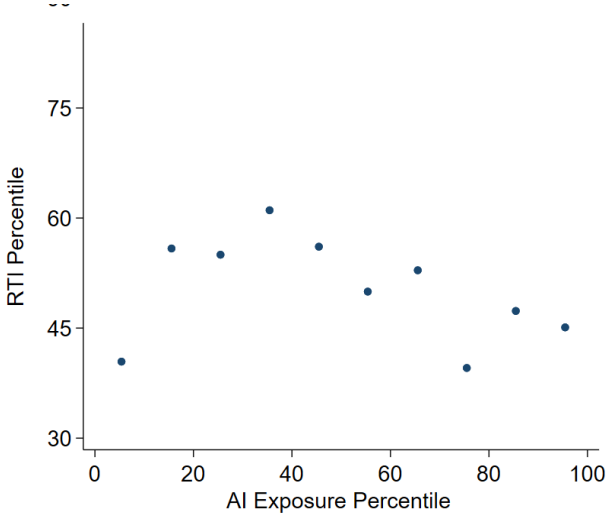

E: Offshorability Exposure

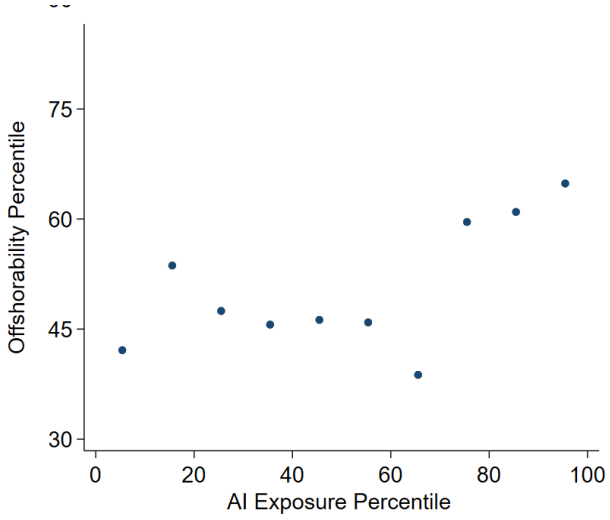

F: Work-From-Home (WFH) Exposure

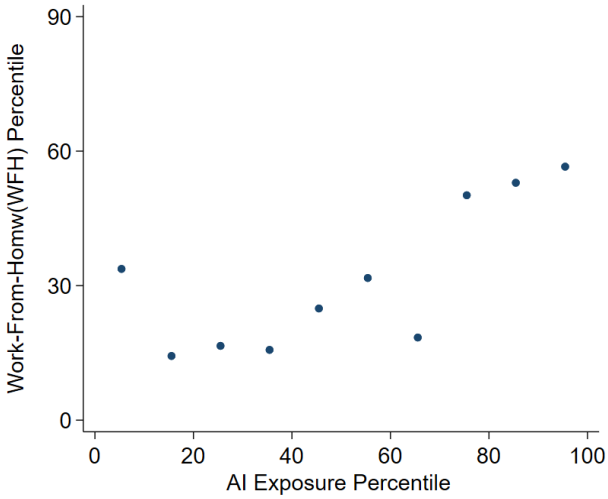

The figure presents the correlation between time-compressed version of general AI exposure (by averaging over the years from 2000–2023) and five time-invariant occupational exposure measures including the AI exposure constructed by Felten et al. (2018), AI exposure and robot exposure constructed by Webb (2019), routine task intensity (RTI) from Autor and Dorn (2013), offshorability potentials from Firpo et al. (2011), and work-from-home (WFH) potentials from Dingel and Neiman (2020). All data series are at the *occ*1990*dd* occupation level. The AI exposure measure is constructed by the authors. Following the literature (e.g., Autor and Dorn, 2013; Webb, 2019), the authors transform all occupation-level exposure scores to percentile ranks and plot the average AI exposure percentile over the other six exposure measures.

Figure OA.3. Google Search Trend of ChatGPT and AI

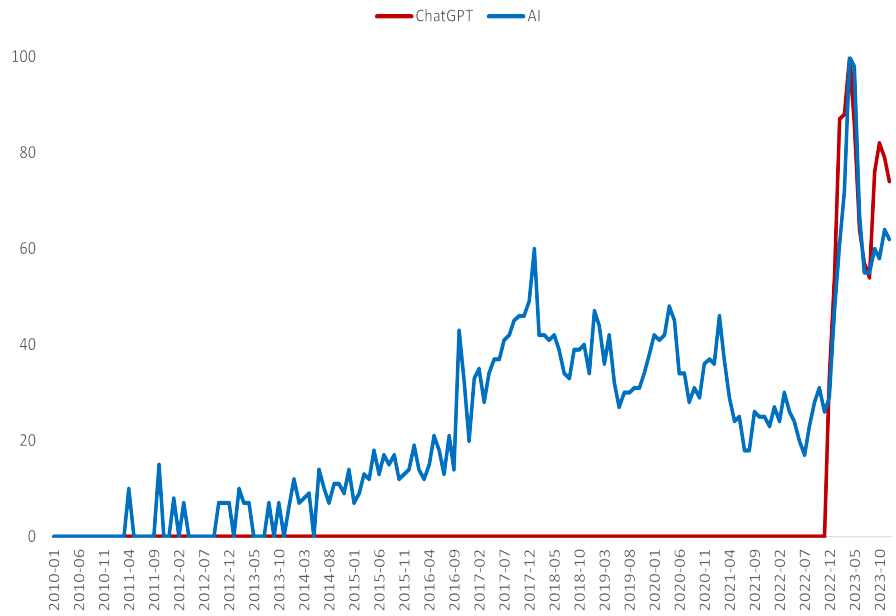

The figure presents the Google search trends of ChatGPT (launched on November 30, 2022) and artificial intelligence (AI) from 2010–2023. The Google Search Trend provides a monthly index scaled from 0 to 100 to indicate the popularity and frequency of particular search terms or topics, where "0" indicates low search volume terms.

D. Tables

Table OA.1: O*NET Database Annula Panel

This table lists the O*NET data release the authors use to construct the annual panel of occupations' job tasks from 2000–2023.

Database	Date Published
O*NET 3.0	8/1/2000
O*NET 3.1	6/1/2001
O*NET 4.0	6/1/2002
O*NET 5.0	4/1/2003
O*NET 6.0	7/1/2004
O*NET 8.0	6/1/2005
O*NET 10.0	6/1/2006
O*NET 12.0	6/1/2007
O*NET 13.0	6/1/2008
O*NET 14.0	6/1/2009
O*NET 15.0	7/1/2010
O*NET 16.0	7/1/2011
O*NET 17.0	7/1/2012
O*NET 18.0	7/1/2013
O*NET 19.0	7/1/2014
O*NET 20.0	8/1/2015
O*NET 21.0	8/1/2016
O*NET 22.0	8/1/2017
O*NET 23.0	8/1/2018
O*NET 24.0	8/1/2019
O*NET 25.0	8/1/2020
O*NET 26.0	8/1/2021
O*NET 27.0	8/1/2022
O*NET 28.0	8/1/2023

Table OA.2: Top Occupations by AI Exposure Scores and and AI Net Complementarity

This table presents occupations with the highest AI exposure (AI^{EXP}) and AI net complementarity (AI^{COMP}), grouped at the 6-digit SOC level for 2023. Occupations are categorized into three groups: high AI^{EXP} & high AI^{COMP}, high AI^{EXP} & low AI^{COMP}, and low AI^{EXP} & low AI^{COMP}. AI^{EXP} and AI^{COMP} is measured by the annual AI-related patent filings from 2018 to 2023, representing the level of AI integration in each occupation.

Occupation Title	O*NET Code	*occ*1990*dd* Title	*occ*1990*dd* Code	AI^{EXP} Score	AI^{EXP} Pct.	AI^{COMP} Score	AI^{COMP} Pct.
High AI^{EXP} & High AI^{COMP}							
Computer and Information Systems Managers	11-3021	Managers and administrators, n.e.c.	22	2.32	100	2.32	100
Electrical Engineers	17-2071	Electrical engineers	55	2.23	100	2.20	100
Computer Hardware Engineers	17-2061	Electrical engineers	55	2.21	100	2.19	100
Inspectors, Testers, Sorters, ...	51-9061	Production checkers, ...	799	2.21	100	1.97	99
Remote Sensing Scientists and Technologists	19-2099	Physical scientists, n.e.c.	76	2.18	100	2.16	100
Operations Research Analysts	15-2031	Operations and systems researchers ...	65	2.14	100	2.12	100
Management Analysts	13-1111	Management analysts	26	2.10	100	2.07	100
Radio Frequency Identification ...	17-2072	Electrical engineers	55	2.10	99	2.06	100
Cartographers and Photogrammetrists	17-1021	Surveyors, cartographers,...	218	2.02	99	1.90	99
Bioinformatics Technicians	43-9111	Statistical clerks	386	2.02	99	1.94	99
High AI^{EXP} & Low AI^{COMP}							
Data Entry Keyers	43-9021	Data entry keyers	385	1.82	96	0.33	23
Log Graders and Scalers	45-4023	Timber, logging, ...	496	1.32	65	-0.34	2
Extruding, Forming, Pressing...	51-9041	Extruding and forming machine ...	755	1.48	80	0.31	21
Office Machine Operators,...	43-9071	Office machine operators, n.e.c.	347	1.48	79	0.33	23
Tellers	43-3071	Bank tellers	383	1.52	83	0.44	29
Transportation Security Screeners	33-9093	Production checkers, ...	36	1.44	77	0.33	23
Parts Salespersons	41-2022	Parts salesperson	275	1.48	80	0.49	31
Rolling Machine Setters, ...	51-4023	Rollers, roll hands, ...	707	1.34	66	0.28	19
Bill and Account Collectors	43-3011	Bill and account collectors	378	1.58	86	0.68	43
Meter Readers, Utilities	43-5041	Meter readers	366	1.40	73	0.50	32
Low AI^{EXP} & Low AI^{COMP}							
Naturopathic Physicians	29-1199	Other health and therapy...	89	0.53	1	0.10	10
Retail Loss Prevention Specialists	33-9099	Protective service, n.e.c.	427	0.53	1	0.06	8
Barbers	39-5011	Barbers	457	0.57	2	-0.03	5
Excavating and Loading Machine ...	53-7032	Excavating and loading machine ...	853	0.59	2	-0.02	6
Welders, Cutters, and ...	51-4121	Welders, solderers, and ...	783	0.59	3	-0.09	3
Shampooers	39-5093	Hairdressers and cosmetologists	458	0.67	5	-0.40	1
Janitors and Cleaners,...	37-2011	Janitors	453	0.70	6	-0.11	3
Sewers, Hand	51-6051	Tailors, dressmakers, and sewers	666	0.71	7	-0.03	5
Dancers	27-2031	Dancers	193	0.71	7	-0.06	5
Slaughterers and Meat Packers	51-3023	Butchers and meat cutters	686	0.74	9	-0.64	1

Table OA.3: AI Exposure and Alternative Specifications

The table reports regression results from alternative specifications estimating the impact of occupational AI exposure on work hours. Detailed information is provided on the following page.

Dep Var	Weekly Work Hours$_{i,o,t}$								
Sample	Full Sample				Exclude				Hourly
		Modified Work			Unemployed		Weekend	Absence	Workers
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
$AI^{EXP}_{o,t-1}$ - TF-IDF	0.049***								
	(2.81)								
$AI^{EXP}_{o,t-1}$		0.024**	0.026***	0.024**	0.020**	0.024***	0.022**	0.028**	0.033**
		(2.46)	(2.94)	(2.46)	(2.59)	(3.17)	(2.39)	(2.62)	(2.29)
$\times\ I(WorkHours - Q4)_{o,t-1}$				-0.017					
				(-1.42)					
$I(Part-time)$			-14.535***			-17.680***			
			(-20.31)			(-34.21)			
Individual characteristics	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Occupational FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
State×Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry×Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year-Month FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Day-of-Week FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	124,059	124,059	124,059	115,407	118,732	118,732	118,918	62,043	58,854
R^2	0.285	0.288	0.313	0.285	0.308	0.351	0.300	0.137	0.284
Adjusted R^2	0.270	0.274	0.299	0.270	0.294	0.337	0.285	0.102	0.253

This table replicates the analyses in Table 3 under alternative model specifications. The dependent variables are the weekly hours spent on market work. The main explanatory variable is occupational AI exposure in percentile rank ($AI_{o,t}^{EXP}$), constructed from job task descriptions and AI patents over a five-year rolling window, as detailed in Section 3.6. We control individual characteristics, including age, the number of children, and a series of indicator variables for gender, educational attainment, marital status, and race. A battery of fixed effects at the following levels are included: occupation, state × year, industry × year, year-month and day-of-week. Standard errors are double clustered by occupation and state. Asterisks denote the significance levels (***=1%, **=5%, *=10%). The model specification for each column is as follows.

(a.) Columns (1): The alternative $AI_{o,t}^{EXP}$ measure is the percentile rank of the TF-IDF–based similarity scores constructed using job task descriptions and AI patents granted in a five-year window ending in the current year. Section B.2.3 of Online Appendix provides detailed descriptions.

(b.) Column (2): The dependent variable is modified market work hours, including hours for commute, work-related travels and social&leisure activities at work.

(c.) Columns (3): The regression additionally controls for an indicator variable for part-time workers.

(d.) Columns (4): Interact AI exposure with a dummy equal to one for occupations whose lagged usual work hours per week are in the top quartile and zero otherwise.

(e.) Columns (5): Currently unemployed individuals are excluded from the sample.

(f.) Columns (6): Currently unemployed individuals are excluded from the sample and the model additionally controls for an indicator variable for part-time workers.

(g.) Columns (7): Individuals who are currently employed but are absent from work on the ATUS interview date are excluded from the sample.

(h.) Columns (8): Individuals surveyed on weekends are excluded from the sample.

(i.) Columns (9): Only individuals compensated on the hourly basis are included in the sample.

Table OA.4: AI Exposure and Workday: Dispersion

The table reports the weighted linear regression results that examine the effect of AI on the within-occupation dispersion of workdays at the occupation × year level derived from the ATUS data from 2004-2023. The regression is weighted by ATUS sample weights at the occupation × year level. The occupation classification is the same as in Table 3. The dependent variable is the standard deviation of weekly hours spent on market work in column (1)-(4) and leisure in column (5)-(8). The main explanatory variables are AI exposure measures in percentile rank at the occupation-year level, calculated from AI-related patents granted over five years ending in the current year (detailed description in Section 3.6). Specifically, they refer to general AI exposure (AI^{EXP}) in columns (1), (2), (5), and (6) and AI net complementarity exposure (AI^{COMP}) in columns (3), (4), (7), and (8). All specifications control for occupation characteristics, including the average age, number of children, and educational attainment, the share of female and married respondents, and year fixed effects. The even columns additionally include occupation fixed effects. Standard errors are clustered by occupation. Asterisks denote significance levels (***=1%, **=5%, *=10%).

DV	Standard Deviation of Weekly Hours$_{o,t}$							
	Work				Leisure			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$AI^{EXP}_{o,t-1}$	-0.001	-0.005*			-0.004	0.000		
	(-0.13)	(-1.76)			(-1.24)	(0.00)		
$AI^{COMP}_{o,t-1}$			0.003	-0.003			-0.006	0.001
			(0.37)	(-0.89)			(-1.20)	(0.16)
Occupation Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Occupation FE	No	Yes	No	Yes	No	Yes	No	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	4,664	4,651	4,664	4,651	4,664	4,651	4,664	4,651
R^2	0.050	0.292	0.050	0.292	0.037	0.178	0.037	0.178
Adjusted R^2	0.045	0.239	0.045	0.238	0.032	0.115	0.032	0.115

Table OA.5: AI Exposure and Workday: Household Allocation

The table reports the weighted linear regression results that examine the heterogeneity effect of AI on work hours sorted based on a subsample of ATUS respondents who report the employment status of their spouse. The ATUS survey sample and occupation classification are the same as in Table 3. The dependent variable is weekly work hours. The first two columns present the results for the full sample of respondents who report their spouses' employment status, and the last two columns present the results for subsamples of those who additionally report their spouses' work hours. The main explanatory variable is AI exposure in percentile rank at the occupation-year level, calculated from AI-related patents granted over the five years ending in the current year (detailed description in Section 3.6). *I(Spouse Employed)* is an indicator that equals one if a respondent's spouse is employed in a given year and zero otherwise. All specifications control for individual characteristics, including age, the number of children, and a series of indicator variables for gender, educational attainment, marital status, and race, and fixed effects at the following levels: state × year, industry × year, year-month and day-of-week. Columns (3) and (4) additionally include occupation fixed effects. Standard errors are double clustered by occupation and state. Asterisks denote significance levels (***=1%, **=5%, *=10%).

Dep Var	Weekly Work Hours$_{i,o,t}$			
			Spouse Work Hours	
			Top 25%	Bottom 75%
	(1)	(2)	(3)	(4)
$AI^{EXP}_{o,t-1}$	0.025***	0.025***	-0.044*	0.039***
	(4.10)	(4.16)	(-1.91)	(3.06)
I(Spouse Employed)$_{i,t}$		-0.704**		
		(-2.29)		
Individual characteristics	Yes	Yes	Yes	Yes
Occupation FE	No	No	Yes	Yes
State × Year FE	Yes	Yes	Yes	Yes
Industry × Year FE	Yes	Yes	Yes	Yes
Year-Month FE	Yes	Yes	Yes	Yes
Day-of-Week FE	Yes	Yes	Yes	Yes
Observations	73,714	73,714	11,651	43,212
R^2	0.336	0.336	0.496	0.367
Adjusted R^2	0.313	0.313	0.386	0.329

Table OA.6: AI Exposure and Alternative Activities

The table reports the weighted linear regressions that examine the effect of occupational AI exposure on time allocated to activities other than market work and leisure at the individual level. The ATUS survey sample and occupation classification are the same as in Table 3. The dependent variable is the weekly hours spent on home production in column (1), child care in column (2), personal education in column (3), job search in column (4), own medical care in column (5), and civic activities in column (6). The main explanatory variable, AI^{EXP}, represents AI exposure in percentile rank at the occupation-year level, calculated from AI-related patents granted over five years ending in the current year (detailed description in Section 3.6). All model specifications control for individual characteristics—age, number of children, gender, educational attainment, marital status, and race—and include fixed effects for occupation, state × year, industry × year, year-month, and day of week. Asterisks denote the significance levels (***=1%, **=5%, *=10%).

Dep Var	Weekly Hours$_{i,o,t}$					
	Home Production	Child Care	Education	Job Search	Own Medical Care	Civic Activities
	(1)	(2)	(3)	(4)	(5)	(6)
$AI^{EXP}_{o,t-1}$	-0.005	0.004	-0.001	-0.000	-0.001	0.004***
	(-0.88)	(0.98)	(-0.49)	(-1.07)	(-1.66)	(2.93)
Individual Characteristics	Yes	Yes	Yes	Yes	Yes	Yes
Occupational FE	Yes	Yes	Yes	Yes	Yes	Yes
State × Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Industry × Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Year × Month FE	Yes	Yes	Yes	Yes	Yes	Yes
Day-of-Week FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	124,059	124,059	124,059	124,059	124,059	124,059
R^2	0.141	0.179	0.121	0.048	0.042	0.080
Adjusted R^2	0.124	0.162	0.103	0.028	0.023	0.061

Table OA.7: AI Exposure and Employment

The table reports the linear regression results that examine the effect of AI on employment at the occupation-year level from 2004-2023. The employment data is provided by Occupational Employment and Wage Statistics (OEWS) from the Bureau of Labor Statistics (BLS). The occupation classification is the same as in Table 3. The dependent variable is the year-over-year change in the natural logarithm of employment. The main explanatory variable is AI exposure in percentile rank at the occupation-year level, calculated from AI-related patents granted over the five years ending in the current year (detailed description in Section 3.6). Specifically, the main explanatory variables are general AI exposure (AI^{EXP}) in column (1), AI net complementarity exposure (AI^{COMP}) in column (2), and changes in the two AI exposure measures in columns (3) and (4), respectively. All models incorporate occupation and year fixed effects. Standard errors are clustered by occupation. Regressions are weighted by lagged employment. Asterisks denote significance levels (***=1%, **=5%, *=10%).

Dep Var	$\Delta\text{Log(Emp)}_{o,t}\times 100$			
	(1)	(2)	(3)	(4)
$AI^{EXP}_{o,t-1}$	0.014			
	(0.36)			
$AI^{COMP}_{o,t-1}$		0.057		
		(0.87)		
$\Delta AI^{EXP}_{o,t}$			-0.072	
			(-1.50)	
$\Delta AI^{COMP}_{o,t}$				-0.037
				(-0.58)
Occupation FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Observations	6,066	6,066	6,066	6,066
R^2	0.144	0.145	0.145	0.094
Adjusted R^2	0.093	0.093	0.093	0.091

Table OA.8: Top Occupations by AI Surveillance

This table presents the top occupations grouped by general AI Surveillance exposure (AI^{SUR}) at the SOC 6-digit level. The procedures measuring AI^{SUR} and its components including $AI^{SUR} - Direction$, $AI^{SUR} - Evaluation$, $AI^{SUR} - Discipline$, are described in Section B.4 of Online Appendix.

Occupation Title	O*NET Code	*occ*1990*dd* Title	Code	AI^{SUR} Score	AI^{SUR} Pct.	AI^{SUR} Pct.		
						Direc -tion	Evalua -tion	Discip -line
		Highest						
Travel agents	41-3041	Transportation ticket ...	318	1.00	100	100	100	100
Air traffic controllers	53-2021	Air traffic controllers	227	0.91	100	100	73	100
Credit analysts	13-2041	Other financial specialists	25	0.91	100	100	100	100
First-line supervisors of retail sales workers	41-1011	Sales supervisors ...	243	0.90	100	100	93	100
Medical records specialists	29-2072	Health record technologists ...	205	0.88	100	94	100	96
Power distributors and dispatchers	51-8012	Power plant operators	695	0.87	99	99	69	100
Bookkeeping, accounting, and auditing clerks	43-3031	Bookkeepers and accounting...	337	0.87	99	95	100	95
Insurance underwriters	13-2053	Insurance underwriters	24	0.86	99	99	100	98
Statistical assistants	43-9111	Statistical clerks	386	0.86	99	83	100	94
Dispatchers, except police, ...	43-5032	Dispatchers	359	0.85	98	99	58	98
		Lowest						
Dancers	27-2031	Dancers	193	0.00	1	1	1	1
Oral and maxillofacial surgeons	29-1022	Dentists	85	0.00	1	1	1	1
Funeral attendants	39-4021	Personal service occupations, n.e.c	469	0.00	1	1	1	1
Automotive glass installers and repairers	49-3022	Auto body repairers	514	0.00	1	3	1	1
Musicians and singers	27-2042	Musicians and composers	186	0.03	3	2	2	1
Dental laboratory technicians	51-9081	Health technologists ...	678	0.06	3	1	3	21
Geographers	19-3092	Social scientists ...	169	0.08	5	6	5	7
Fine artists, including painters, sculptors...	27-1013	Painters, sculptors, ...	188	0.09	6	7	6	6
Clergy	21-2011	Religious workers, n.e.c.	176	0.10	6	7	6	5
Optometrists	29-1041	Optometrists	87	0.10	6	1	7	37

Table OA.8: Exposure to AI Surveillance Technology and Workday: Self-Employed

The table reports the weighted linear regression results based on individual responses of self-employed working remotely in the ATUS survey from 2015 to 2023. The occupation classification is the same as in Table 3. Remote workers are defined as those in occupations with a work-from-home (WFH) feasibility index from Dingel and Neiman (2020) equals one. The dependent variable is weekly work hours. The main explanatory variable, AI_o^{SUR}, is AI surveillance exposure at the occupation level (detailed description in Section B.4 of Online Appendix) and transformed to percentile ranks (e.g., Autor and Dorn, 2013; Webb, 2019). Specifically, it refers to general AI surveillance exposure in column (1) and decomposed AI surveillance exposure in column (2)-(4) as specified in the third row. *POST* dummy equals one for the years since 2020. All specifications control for individual characteristics, including age, the number of children, and a series of indicator variables for gender, educational attainment, marital status, and race, and fixed effects at the following levels: occupation, state × year, industry × year, year-month, and day-of-week. Standard errors are double clustered by occupation and state. Asterisks denote significance levels (***=1%, **=5%, *=10%).

Sample	Weekly Work Hours$_{i,o,t}$			
Indep Var	Overall	Direction	Evaluation	Discipline
	(1)	(2)	(3)	(4)
$AI_o^{SUR} \times POST_t$	0.010	0.061	0.006	0.032
	(0.05)	(0.29)	(0.03)	(0.16)
Individual characteristics	Yes	Yes	Yes	Yes
Occupational FE	Yes	Yes	Yes	Yes
State × Year FE	Yes	Yes	Yes	Yes
Industry × Year FE	Yes	Yes	Yes	Yes
Year-Month FE	Yes	Yes	Yes	Yes
Day-of-Week FE	Yes	Yes	Yes	Yes
Observations	644	644	644	644
R^2	0.839	0.839	0.839	0.839
Adjusted R^2	0.497	0.497	0.497	0.497

Table OA.8: AI Exposure and Employee Satisfaction: Introduction of ChatGPT

The table presents estimates from linear regressions examining the effects of occupational AI exposure on employee ratings at the occupation (o) × firm (i) × year (t) level around the introduction of ChatGPT. The occupation classification is the same as in Table 3. The sample covers data from private and public firms in the Glassdoor database from 2022 to June 2023. The dependent variables are: 100 times the overall satisfaction rating in columns (1)–(2) and the Work-Life Balance (WLB) ratings in columns (3)–(4). The main explanatory variable represents AI exposure measures at the occupation in 2021, expressed in percentile ranks, and is based on AI-related patents granted in a five-year window ending in the current year (detailed description in Section 3.6). *POST* dummy equals one for year 2023. All specifications include lagged employee review counts, the average seniority and remote work index of the occupation-by-firm cohort from Revelio, and occupation fixed effects. The odd columns additionally include firm and year fixed effects, whereas the even columns include firm-by-year fixed effects. Standard errors are clustered by occupation. Asterisks denote significance levels (***=1%, **=5%, *=10%).

Dep Var	$\text{Rating}_{o,i,t}$			
	Overall		WLB	
	(1)	(2)	(3)	(4)
$AI^{EXP}_{o,2021} \times POST_t$	-0.062**	-0.054**	-0.070**	-0.061**
	(-2.50)	(-2.11)	(-2.55)	(-2.26)
Cohort Controls	Yes	Yes	Yes	Yes
Occupation FE	Yes	Yes	Yes	Yes
Firm FE	Yes	No	Yes	No
Year FE	Yes	No	Yes	No
Firm×Year FE	No	Yes	No	Yes
Observations	64,813	64,813	59,484	59,482
R^2	0.200	0.228	0.215	0.243
Adjusted R^2	0.164	0.171	0.176	0.182

Table OA.9: AI Exposure and Workday: In Relation to Competition

The table reports the weighted linear regression results that estimate the heterogeneity effect of AI on work hours sorted by labor market and product market competition. The ATUS sample and occupation classification are the same as in Table 3. The dependent variable is the weekly hours spent on market work in columns (1)-(2) and leisure in columns (3)-(4). The main explanatory variable, AI^{EXP}, is AI exposure in percentile rank at the occupation-year level, calculated from AI-related patents granted over five years ending in the current year (detailed description in Section 3.6). Two proxies for firms' labor market competition at the state-occupation level are specified: the employment concentration across firms measured by the Herfindahl-Hirschman Index (HHI) in columns (1) and firms' talent retention pressure (TRP) in columns (2). The state-occupation-level employment HHI is calculated using the employment and state information of public firms from LinkedIn. TRP is calculated as the job vacancy-to-employment ratio (V/E) at the state-occupation level, scaled by each occupation's cognitive skill scores, using job vacancy data from Burning Glass, employment data from the OEWS, and cognitive skill scores from O*NET. *I(Worker Power vs. Firm)* is the indicator of workers' bargaining power over firms that equals one if the measure for firms' labor market competition is in the top quartile (i.e., the employment HHI in the bottom quartile or TRP in the top quartile), and zero otherwise. Two proxies represent firms' product market power at the industry-level: product similarity in column (3) and product market concentration HHI in column (4), derived from firm-level scores from Hoberg and Phillips (2016) weighted by Compustat sales. *I(Firm Power vs. Consumer)* is an indicator of firms' product market power relative to consumers, which equals one if the product similarity is in the bottom quartile or the product HHI is in the top quartile, and zero otherwise. All specifications incorporate individual-level controls, including age, the number of children, and a series of indicator variables for gender, educational attainment, marital status, and race, and fixed effects at the following levels: occupation, state × year, industry × year, year-month and day-of-week. Standard errors are double clustered by occupation and state. Asterisks denote significance levels (***=1%, **=5%, *=10%).

Dep Var	Weekly Work Hours$_{i,o,t}$			
	Labor Market Competition		Product Market Competition	
Factor	HHI	Talent	Similarity	HHI
	(1)	(2)	(3)	(4)
$AI^{EXP}_{o,t-1}$	0.025**	0.065***	0.022**	0.027***
	(2.50)	(4.06)	(2.35)	(2.92)
× *I(Worker Power vs. Firm)*$_{i,s,t-1}$	-0.011	-0.041*		
	(-1.28)	(-1.83)		
× *I(Firm Power vs. Consumer)*$_{j,t-1}$			-0.004	-0.027**
			(-0.39)	(-2.56)
Individual Characteristics	Yes	Yes	Yes	Yes
Occupational FE	Yes	Yes	Yes	Yes
State× Year FE	Yes	Yes	Yes	Yes
Industry× Year FE	Yes	Yes	Yes	Yes
Year× Month FE	Yes	Yes	Yes	Yes
Day-of-Week FE	Yes	Yes	Yes	Yes
Observations	102,434	48,349	114,661	114,661
R^2	0.289	0.290	0.283	0.283
Adjusted R^2	0.272	0.272	0.268	0.268

Table OA.10: Example Employee Reviews with Complaints

This table presents example employee reviews with negative sentiment toward surveillance and employment risk. Each review includes a "Pros" section, highlighting positive remarks about the firm, and a "Cons" section with negative remarks.

Occupation	Firm	Review Date	Review Text
		Complaints - Surveillance	
51-3011 Bakers	Target Corp.	2016-12-01	**Pros**: Hours are given when holidays come around and activities for employees like bbq, food, and other activities. Pay is a little bit abovd minimum wage. **Cons**: Lack of help and always being watched by hidden **cameras**. Management is very poor and to be working under staff it really is difficult to keep up. Redcards are the new hard working employee (creditcards).
53-6021 Parking Attendants	Amazon.com, Inc.	2020-10-27	**Pros**: pretty lay back job, lots of retired worker **Cons**: **spying** on you using **camera** and AI
43-4181 Reservation and Transportation Ticket Agents and Travel Clerks	United Airlines, Inc.	2023-03-13	**Pros**: Flight benefits were my reason to sign on with the company. The entry-level pay is comparable to that of other airlines with a slight differential if you speak a foreign language. work-life-balance can be good as long as there are no flight disruptions, otherwise expected to work in mandatory overtime. Flexibility can be high if you're able to trade away your shifts. **Cons**: Micro-management by immediate supervisors and managers is at a very high level. Employee **monitoring** can be very stressful and lead to extreme anxiety. Company focus on performance metrics is almost inhumane. Employees are numbers and high replaceable. Toxic culture of high expectation and intimidation.

Complaints - Employment Risk			
11-2022 Sales Managers	Amazon.com, Inc.	2022-09-21	**Pros**: If you have a great manager, the work-life-balance is good. Makes the job fun. **Cons**: If you have a bad manager, which there are many, the team will have issues, managers are **firing** more than they are teaching and the company is a nightmare to work for. All of the good employees leave and the crap ones stay behind.
15-1252 Software Developers	Google LLC	2023-01-24	**Pros**: Google used to be a great place to get away from office politicking and just focus on doing the work you love. **Cons**: Recent **layoffs** were incredibly demoralizing. With **layoff** decisions being made at a level so far removed from the people doing the work that keeps things running, basically nobody's job is safe.
27-2012 Producers and Directors	Meta Platforms, Inc.	2023-02-27	**Pros**: Overall, love the people I work with, love Meta's dedication to DEI in all we do, and we have an interesting value proposition for the future. **Cons**: Huge lack of planning org-wide. This leads to continous **layoffs/uncertainty** and lots of duplication of work/inefficient processes. Also still fighting a culture that has been bottom up for so long, which in terms of being impactful and getting things done that benefit the business, not just individuals, makes things difficult as a manager.