Job Transformation, Specialization, and the Labor Market Effects of AI*

Lukas B. Freund Lukas F. Mann

This version: November 11, 2025
First version: August 16, 2025

Please click here for the latest version.

Abstract

Who will gain and who will lose as AI automates tasks? While much of the discourse focuses on job displacement, we show that *job transformation*—a shift in the task content of jobs—creates large and heterogeneous earnings effects. We develop a quantitative, task-based model where occupations bundle multiple tasks and workers with heterogeneous portfolios of task-specific skills select into occupations by comparative advantage. Automation shifts the relative importance of tasks within each occupation, inducing wage effects that we characterize analytically. To quantify these effects, we measure the task content of jobs using natural language processing and estimate the distribution of task-specific skills. We construct projections of automation effects due to large language models (LLMs), exploiting a mapping between model tasks and automation exposure measures. Within highly exposed occupations, like office and administrative roles, workers specialized in information-processing tasks leave and suffer wage losses. By contrast, those specialized in customer-facing and coordination tasks stay and experience wage gains as work rebalances toward their strengths. Our findings challenge the common assumption that occupational automation exposure necessarily implies individual wage losses; and highlight that AI, through job transformation, may be disruptive even absent job displacement.

*Freund: Boston College, lukas.beat.freund@gmail.com, www.lukasfreund.com; Mann: Arizona State University, lfmann@asu.edu, www.lukasmann.com.

We thank David Autor, Maarten De Ridder, John Grigsby, Jonathan Heathcote, Kyle Herkenhoff, Anders Humlum, Huiyu Li, Ilse Lindenlaub (discussant), Ellen McGrattan, Simon Mongey, Sergio Ocampo Díaz, Pascual Restrepo, Benjamin Schoefer, Ayşegül Şahin, Todd Schoellman, Joe Stiglitz, Mike Webb and Nathan Zorzi for helpful conversations and comments. We also thank seminar and conference audiences at Columbia, the Minneapolis Fed, Boston University, Harvard Kennedy School, Harvard, as well as at Midwest Macro 2025, SED 2025, BSE Summer Forum 2025, SITE 2025 and the NBER EFG Fall Meeting 2025 for useful comments and suggestions. Alli Tucker, ChatGPT and Claude provided outstanding research assistance. A conference draft of this paper was circulated with the title "For Whom the Bot Tolls: Specialization and the Earnings Effects of AI."

1 Introduction

Rapid advances in artificial intelligence (AI) raise the prospect of machines taking over an expanding set of tasks. While the discourse about the labor market consequences of such a wave of automation often centers on entire jobs being eliminated (Frey and Osborne, 2017; Susskind, 2020), history suggests that automation transforms what tasks workers perform within jobs long before it erases them (Autor *et al.*, 2003). We refer to this process as *job transformation*. During the industrial revolution, weavers continued working in large numbers after power-looms were introduced, but their responsibilities shifted toward fixing mechanical issues and coordinating workflow across multiple looms (Bessen, 2012). In the late 20th century, CNC tools shifted machinists' roles from routine tasks like positioning tools to programming, monitoring and correcting digital processes (Bartel *et al.*, 2007).[1] Recent studies suggest such task shifts will likely play a major role in the case of AI, too (Bonney *et al.*, 2024; Gathmann *et al.*, 2024).[2] In this paper we ask: which workers will win and which will lose from AI-induced job transformation?

While job transformation appears of first-order importance for the labor market effects of AI, quantifying its role presents significant measurement challenges. First, the sign and magnitude of job transformation effects depend on workers' entire portfolios of task-specific skills, which are typically unobserved. Second, analysis requires knowing which specific tasks are automated. State-of-the-art models of automation therefore typically abstract from job transformation, which allows characterizing wage effects through the competing forces of labor share declines and productivity gains (Acemoglu and Restrepo, 2018a,b).

We develop a task-based model of job transformation and use it to construct a quantitative scenario of how automation by large language models (LLMs) will affect wages for heterogeneously skilled workers. To overcome the two measurement challenges, we map the model to granular tasks linked to prominent automation exposure measures (Webb, 2019; Eloundou *et al.*, 2023) and exploit the model's structure to estimate the distribution of task-specific skills.

Our analysis yields three main findings. First, as LLM-driven automation of information-processing tasks generates larger shifts in occupational skill composition than past automation by industrial robots, occupational averages provide limited guidance for worker-level outcomes. Second, workers specialized in LLM-exposed tasks leave their transformed jobs and experience

[1] Appendix B.1 provides more details. Systematic studies (e.g., Spitz-Oener, 2006; Atalay *et al.*, 2020) corroborate these vignettes, demonstrating that the shift from routine tasks toward non-routine analytic and social tasks over recent decades arose predominantly from changes in task content within occupations rather than shifts in occupations' employment shares.

[2] Reporting in the Financial Times ("Disrupted or displaced? How AI is shaking up jobs," June 08 2025) provides complementary anecdotal evidence. For example, "According to PwC, the mix of capabilities sought by employers is changing 66 per cent faster in occupations most exposed to AI, such as financial analysts, than in those least exposed, such as physical therapists."

wage losses. Third, the same shock creates two groups of winners: incumbent workers remaining in highly exposed office jobs, freed to spend more time on customer-facing and coordination tasks in which they excel; and workers previously deterred from highly exposed occupations by skill barriers in now-automated tasks, who following automation switch into these transformed roles. Overall, our findings challenge the common practice of equating occupational automation exposure with individual wage losses; and highlight that automation, through job transformation, can create large and heterogeneous wage effects even absent job elimination.

Theory. We introduce a model of automation-driven job transformation, building on the canonical task-based theory.[3] In our model, task aggregation occurs at the occupational level, with different occupations attaching heterogeneous weights to tasks. Each worker has a portfolio of task-specific skills. Tasks are assigned to humans or machines. We integrate this production framework with Roy (1951)-style occupational choice based on comparative advantage. To expose which automation wage effects arise solely from job transformation, our model is deliberately partial equilibrium, treating occupational output prices as fixed.

A key ingredient of our model of automation is *task bundling*: to produce output, a worker must perform all tasks in an occupation, including tasks they are relatively less skilled at. We do not model the sources of bundling but view bundling as eminently plausible—most economists may be more productive at data analysis or math than at emailing, yet emailing remains part of their job. In our empirical analysis, we discipline the degree of bundling by measuring the task content of jobs. Task bundling significantly affects how automation works. Automation reassigns tasks from labor to machines, and in the presence of bundling, this creates job transformation: In occupations involving the automated task, workers spend less time on that task—none if it is fully automated—while time on all other bundled tasks increases proportionately.

The model provides an intuitive characterization of wages and how their response to automation are shaped by task bundling and worker specialization. Wage levels reflect both absolute advantage (i.e., average skill) and comparative advantage (i.e., alignment between skill specialization and occupational task requirements). Following automation, individual wages change not only due to the standard displacement and productivity effects, whose net effect appears as an occupation-level shifter, but also due to job transformation: Workers relatively more skilled at automated tasks than at non-automated tasks tend to lose, whereas gains accrue to those relatively unskilled at automated tasks. Consider lawyers as an example. According to the model, when document processing is automated, a lawyer specialized in this task will lose, whereas a colleague with a comparative advantage in client engagement will gain.

We derive an analytical decomposition of job transformation effects for average occupational wages: First, they depend on whether the occupation's task content shifts toward tasks for which

[3] See, in particular, Autor *et al.* (2003); Acemoglu and Autor (2011); Acemoglu and Restrepo (2018b).

the average worker is more skilled (*task upgrading*) or less skilled (*task downgrading*). Second, they tend to be more negative for more exposed occupations, which often attract workers who are exceptionally skilled at the automated task (*selection*). Third, they tend to rise when automation leads to large shifts in the skill composition of an occupation (*re-sorting*). The absolute and relative magnitude of these job transformation effects depend on the underlying distribution of multi-dimensional skills, which is not directly observed.

Measuring the skill distribution. To make progress, we leverage our model's structure to estimate the distribution of multi-dimensional skills using maximum likelihood. Intuitively, the identification works as follows. Consider two occupations, economists and software engineers, where both code, but economists also write. Identification comes from two sources: wage comparisons and occupational choices. If we observe a worker in both occupations, their wage as a software engineer reveals their coding skill, letting us infer their writing skill from their economist wage. If we only observe the worker as an economist, say, this reveals their coding skills fall below the threshold for choosing software engineering. We formalize these intuitions as a maximum-likelihood estimation problem to recover the means and variance-covariance structure of skills alongside other structural parameters.

We implement this approach relying only on publicly available data, the National Longitudinal Survey of Youth (NLSY) 1979, and by measuring occupational task weights using natural language processing (NLP) techniques. For the latter, we cluster approximately 19,000 detailed occupation-specific tasks from the Occupational Information Network (O*NET) into 38 tractable and interpretable task categories. Our theory implies that occupational task weights reflect optimal time allocation, and we use an LLM to measure these weights as time shares across tasks for each occupation.[4] The resulting task-weight matrix indicates task bundling: only 2% of occupations have a single task comprising more than half their time.

The estimated model fits several salient empirical moments well. Beyond generating realistic wage distributions within and between occupations, the model predicts that, conditional on switching occupations, workers tend to move to jobs with task requirements similar to their origin occupation — consistent with evidence for skills being task-specific but hard to explain if skills were fully general or occupation-specific (Gathmann and Schönberg, 2010).

Mapping to automation exposure measures. Quantifying automation-induced job transformation effects for current or future technologies requires knowing which exact tasks are being, or will be, automated. This is the second of the two aforementioned measurement challenges.[5] Our framework enables us to overcome this challenge: The model tasks have a direct counterpart in

[4] We extensively validate this LLM-based approach, including through comparisons to worker time diary data.

[5] Quantifying job transformation effects thus requires information beyond the total task displacement as measured by the overall labor share decline.

the data in the form of (aggregated) O*NET task categories. As the underlying detailed O*NET tasks are widely used to construct empirical task-level exposure measures, we can leverage this mapping to identify task-specific automation shocks.

For our main analysis, we construct projections of automation effects due to LLMs, given the rapid diffusion of LLMs (Bick *et al.*, 2024) and policymaker interest in their labor market consequences.[6] To identify tasks most exposed to LLMs, we draw on Eloundou *et al.*'s (2023) exposure measures, which reveals several information-processing tasks—common in office and administrative support roles, such as financial clerks and record clerks—as most exposed. To compare effects across technologies, we also evaluate the consequences of industrial robot automation, leveraging the exposure measure developed by Webb (2019).

Results. Our analysis shows that LLM-driven job transformation creates large and heterogeneous worker-level earnings effects. Our first result is that average occupational wage changes are a poor indicator for individual incumbents' outcomes. While average wages in the most exposed occupations increase following LLM automation, this largely reflects compositional changes and, thus, re-sorting effects. These arise because job transformation significantly alters skill requirements: customer-facing and coordination tasks rise in significance as information-processing tasks are automated. More broadly, we show that re-sorting wage effects are larger for shocks affecting tasks with more dispersed skills. Our estimates show that LLM-exposed tasks exhibit more skill dispersion across workers than robot-exposed tasks, suggesting occupation-level comparisons provide worse guidance for the AI era than for historical automation.

Our second result highlights which individual workers lose most from LLM automation: those who chose office and administrative support roles because they excelled at informaton-processing tasks. LLM automation erases these workers' comparative advantage in their current jobs, leading them to exit for the next-best occupation and suffer large wage losses. The experience of these incumbent-leavers reflects the selection effect. This creates a negative relationship between incumbent wages and exposure, contrasting sharply with occupation-level outcomes.

Our third result highlights two groups of winners. For one, not all incumbents lose. Workers who sorted into highly exposed occupations because they excel at tasks bundled with the automated task – such as customer service or administrative coordination – experience wage gains from task upgrading. Automation frees them to spend more time on non-automated tasks where they are relatively more productive. The largest wage gains, however, accrue to workers who switch *into* highly exposed occupations. These in-switchers were previously deterred from

[6] See, for instance, IMF research (Cazzaniga, 2024) and recent speeches by IMF and ECB heads Georgieva and Lagarde: https://www.imf.org/en/Blogs/Articles/2024/01/14/ai-will-transform-the-global-economy-lets-make-sure-it-benefits-humanity and https://www.ecb.europa.eu/press/key/date/2025/html/ecb.sp250401_1~d6c9d8df11.en.html.

exposed occupations by the large weight on information-processing tasks. Automation thus removes a skill-based entry barrier, generating significant wage gains for these workers.

Zooming out, our findings carry two major implications. First, they suggest that the widespread practice of equating occupation-level automation exposure with adverse worker-level impacts can be misleading. Even absent positive productivity effects at the occupation level, workers within any occupation may win or lose depending on their skill specialization. In brief, exposure measures are best interpreted as indicating *potential* for change; they need to be paired with a structural model that carefully maps exposure into wages, at least for forward-looking analyses. Second, an absence of evidence documenting widespread job elimination should not be interpreted to imply that AI lacks major labor market effects. In our model, AI automation generates wage effects by transforming how workers spend their time.

Literature. Our paper contributes to a burgeoning literature evaluating the labor market consequences of AI. One influential strand empirically quantifies task exposure to new technologies, AI foremost among them, drawing on patents (Webb, 2019; Kogan *et al.*, 2023), capability-specific AI benchmarks (Felten *et al.*, 2018, 2021), and expert or machine judgment (Brynjolfsson *et al.*, 2018; Eloundou *et al.*, 2023). These exposure measures alone cannot predict earnings consequences. Our paper complements this work by offering a structural approach to map task exposure measures to individual-level labor market outcomes. Our findings underscore that similarly exposed individuals may experience very different earnings effects.

Methodologically, our work belongs to a second strand that uses structural models. The most closely related paper is Hampole *et al.* (2025).[7] Hampole *et al.* (2025) use CV and job posting data to construct firm- and time-varying measures of exposure to existing machine learning/AI techniques, which allows them to study heterogeneity across firms — whereas our analysis is largely silent on firm heterogeneity. Like ours, their model features occupations comprising multiple tasks. We make two distinct contributions. First, whereas workers in their model are ex-ante identical, our theory features multi-dimensional skill heterogeneity. We contribute a methodology to empirically discipline this heterogeneity. Theory and skill measurement jointly allow us to show how automation, through job transformation, differentially affects worker earnings based on heterogeneous skills. This enables us to show that workers in the same occupation may fare differently depending on their specialization. Second, our model links to forward-looking task exposure measures, enabling us to quantify the labor market consequences of generative AI — central to policymaker and societal concerns.[8]

[7] Several concurrent works are developing quantitative task-based models of AI (Fan and Restrepo, 2025; Lashkari *et al.*, 2025; Chequer *et al.*, 2025; Althoff and Reichardt, 2025). A comparative discussion of these complementary approaches will follow once working papers become available.

[8] Three other strands of AI research merit highlighting. The first comprises surveys characterizing adoption patterns and early trends in work reorganization (Bick *et al.*, 2024; Humlum and Vestergaard, 2025b,a). The second

Second, we contribute to the literature of task-based production.[9] Our first contribution is to integrate this theory of production with occupational choice (Dix-Carneiro, 2014; Hsieh *et al.*, 2019; Traiberman, 2019).[10] Beyond being instrumental for our measurement strategy, this approach captures occupational reallocation as an empirically salient adjustment margin (Dauth *et al.*, 2021; Boustan *et al.*, 2022). Second, and more importantly, we characterize automation effects through job transformation that are typically assumed away by holding tasks fixed.[11] In terms of theory, this is a technically straightforward extension of the canonical task-based model. Our primary contribution lies in measurement: quantifying job transformation effects requires knowledge of the distribution of task-specific skills, which we estimate. Moreover, since the labor share is no longer a sufficient statistic for displacement, we show how to leverage measures of technology-specific task exposure to construct automation shocks. Our emphasis on task bundling relates to the contemporaneous work of Autor and Thompson (2025), though the papers differ substantially in methodology and focus. Autor and Thompson's (2025) model features a strict expertise hierarchy where more expert workers compete with less expert workers but not vice versa. Our model has no occupation-level expertise concept; instead, we estimate the distribution of skills across tasks, so some tasks emerge as "expert"' ex post. Autor and Thompson (2025) use their model with a novel reduced-form approach to resolve the historical puzzle of why routine task automation often raised wages in routine-intensive occupations despite employment declines. We provide a fully structural quantitative analysis of earnings effects from ongoing or future AI automation.

Third, our focus on job transformation is motivated by a large empirical literature. Starting with the seminal work of Autor *et al.* (2003), numerous studies have documented significant shifts in task requirements within jobs over time.[12] However, this work has largely not connected tasks, which characterize jobs, to multi-dimensional skills, which characterize workers and are typically unobservable. This paper provides a structural framework to measure the distribution

studies whether AI can accelerate economic growth rates by automating tasks in R&D (Aghion *et al.*, 2017; Jones, 2022). This mechanism is beyond our scope but would roughly correspond to a level shift in incomes, leaving open questions about the distribution of gains that are at the core of this paper (cf. Autor, 2015, p.28). The third involves RCTs that causally identify the productivity effects of generative AI adoption in narrowly defined contexts (e.g., Noy and Zhang, 2023; Brynjolfsson *et al.*, 2025). Further, Ide and Talamas (2025) uses a knowledge-based hierarchy model à la Garicano (2000) to study the effects of AI on firm organization and earnings.

[9] See, among others, Autor *et al.* (2003); Acemoglu and Autor (2011); Acemoglu and Restrepo (2018b); Ocampo Díaz (2022); Moll *et al.* (2022); Freund (2023); Restrepo (2024).

[10] Several studies focus on occupational reallocation in response to technological shocks, with an emphasis on GE effects (Humlum, 2019; del Rio-Chanona *et al.*, 2021; Bocquet, 2022; Fan, 2025; Böhm *et al.*, 2025). Relatedly, Grigsby and Zorzi (2025) study the distributional effects of the green transition in a model with skill heterogeneity.

[11] An additional difference is that, in our approach, positive productivity effects accrue only to exposed occupations, because automated tasks are not bundled together with every other task as in existing models.

[12] These studies draw on the Dictionary of Occupational Titles (DOT) and O*NET, worker surveys (Autor and Handel, 2013; Spitz-Oener, 2006), and job ads (Atalay *et al.*, 2020). Also see Lin (2011) and Autor *et al.* (2024).

of task-specific skills and to quantify the heterogeneous earnings effects of automation-induced job transformation.

Fourth, we contribute to research demonstrating the importance of multi-dimensional skills for labor markets.[13] Our primary contribution lies in measurement: developing and implementing a method to estimate the distribution of task-specific skills. We thus relate to Guvenen *et al.* (2020), Lise and Postel-Vinay (2020) and Baley *et al.* (2022), who used military test scores to measure skills among workers. Our methodology is not anchored by observed test scores but offers two advantages. First, flexibility: Our methodology can be applied to any large-scale worker dataset with information on occupations and wages, without requiring (rare) data on test scores. Second, we estimate skill distributions in potentially high-dimensional task spaces rather than being restricted to low-dimensional categories like cognitive, manual, and interpersonal skills. This allows us to connect skills to granular tasks, as considered in the literature on automation exposure.[14] Overall, our paper closes the gap between the literatures on multi-dimensional skills – which has thought carefully about skill measurement and sorting but relies on abstract notions of technological change – and task-based production – which highlights how the demand for specific skills is shaped by automation.[15]

Finally, by integrating an LLM into our empirical workflow, we relate to a nascent literature showing how these tools can be leveraged for economics research (e.g., Kogan *et al.*, 2023; Athey *et al.*, 2024). Our use of LLMs to cluster tasks and measure occupational task allocation resembles the data processing and classification use cases discussed by Dell (2024). We show that this approach produces results consistent with established measurement frameworks while offering greater flexibility.

Outline. Section 2 presents the theory, Section 3 takes it to the data, Section 4 evaluates the labor market effects of LLMs, and Section 5 offers a concluding discussion.

[13] In terms of theory, Lindenlaub (2017) likewise studies multidimensional matching between workers and jobs and how technological change shapes it. While Lindenlaub (2017) focuses on shifts in complementarity between skills and production requirements, we adopt a task-based production approach to study automation. Our model also resembles Lazear's (2009) skills-weights approach, treating skills not as inherently specific to a single production unit — firms in Lazear's work, occupations in ours – but recognizing that different units attach heterogeneous weights to different skills. For surveys see Deming (2023) and Woessmann (2024).

[14] Grigsby (2023), while pursuing a different question, likewise infers the multidimensional skill distribution from occupational choices and wages. The most important of several differences in methodology is that in Grigsby's (2023) approach, a task corresponds to a group of occupations, whereas we conceptualize occupations as bundles of tasks and estimate the distribution of these granular skills. This distinction between occupations and tasks is essential for studying the consequences of job transformation.

[15] Woessmann (2024, p.4) summarizes the gap in the literature this paper helps fill: "[Although] worker skills motivate the entire task-based approach to how labor markets adjust to technological change, the consideration of multidimensional tasks has not been matched by multidimensional measurement of skills on the empirical side. While the tasks required in different jobs are richly described, worker skills are still mostly proxied rudimentarily by educational degree."

2 Theoretical Framework

In this section we set out the theoretical environment (Section 2.1), derive optimality conditions and define the equilibrium (Section 2.2). We then define automation in the context of the model and characterize its effects on wages (Section 2.3).

2.1 Environment

Time is discrete and runs forever. The economy is populated by workers and entrepreneurs who produce and consume a single, homogeneous numeraire good.

Workers. There is a unit mass of infinitely lived workers. Before the onset of time, each worker draws and observes their skill vector $s_i \in \mathbb{R}^{n_{\text{skill}}}$, where $s_i \sim \mathcal{N}(\bar{s}, \Sigma_s)$. This skill vector remains fixed forever.[16] In each period t, a worker draws two shocks: a productivity shock $\varepsilon_{i,t} \sim \mathcal{N}(0, \varsigma^2)$, and a vector of occupation-specific preference shocks $u_{i,\cdot,t} \in \mathbb{R}^{n_{\text{occ}}}$, $u_{i,\cdot,t} \sim \text{Gumbel}(0, \nu)$.

Production. Production occurs across n_{occ} occupations indexed by $o \in \mathcal{O}$. Production in occupation o requires that a series of tasks $\tau \in \mathcal{T}$ be carried out; what distinguishes occupations from each other are the weights $\{\alpha_{o,\tau}\}_{\tau \in \mathcal{T}}$ attached to these tasks, with $\sum_{\mathcal{T}} \alpha_{o,\tau} = 1 \, \forall o \in \mathcal{O}$. Concretely, the amount of output in an occupation o job is determined by a Cobb-Douglas aggregator with occupation-specific weights $\alpha_{o,\tau}$. Hence, the output of a worker i in occupation o is

$$y_{i,o,t} = \prod_{\tau \in \mathcal{T}} x_{i,\tau,t}^{\alpha_{o,\tau}} \tag{1}$$

where $x_{i,\tau,t}$ is the amount of task τ used in production.[17] We interpret these tasks as concrete work steps that need to be performed in a given occupation, such as analyzing business data, moving materials, delivering instruction, etc. A task can be produced using (i) the worker's time or (ii) machine capital. Machine capital has a productivity $\exp(z_\tau)$ at task τ and can be rented from an infinitely elastic capital market at exogenous rate r.[18] We denote the set of tasks

[16] A couple of remarks regarding the assumption that skills are time-invariant. In part, the assumption is motivated by computational constraints in the estimation of skills in Section 3. Additionally, in Section 2.3.3, it enables us to derive closed-form analytical approximations that allow us to sharply characterize the effects of job transformation.

[17] The unit elasticity of substitution across bundled tasks implicit in equation (1) represents a common baseline in the literature (e.g. Acemoglu and Restrepo, 2022, pp. 1986) and carries some important advantages. First, it provides a transparent way for measuring $\{\alpha_{o,\tau}\}_{o \in \mathcal{O}}$, as described in Section 3.2, because task shares are invariant to shifts in task-specific skill. Second, it confers significant tractability when estimating the skill distribution by producing a log-linear wage equation. Third, because under Cobb-Douglas a productivity-enhancing automation shock does not mechanically increase the relative demand for human-performed tasks (and hence wages), this assumption transparently isolates wage effects arising from the interaction of shifting task weights and skill specialization.

[18] An infinitely elastic capital supply will tend to raise average wages following the adoption of a new automation technology (Caselli and Manning, 2019), relative to the case of a fixed capital stock (Acemoglu and Restrepo, 2018b,

produced with human labor as $\mathcal{T}_l$ and the set of tasks produced with machine capital as $\mathcal{T}_m$. For now, we treat these sets as exogenous and assume only that they do not depend on the specific occupation nor the skill of any individual worker. For the purposes of formalizing automation, Section 2.3.1 discusses a set of additional assumptions under which $(\mathcal{T}_l, \mathcal{T}_m)$ can be endogenized.

A competitive firm sector sets (log) wages $w_{i,o,t}$ as a function of each worker's skill $s_{i,\tau}$ and idiosyncratic shock $\varepsilon_{i,t}$. The firm freely allocates the worker's unit measure of labor across tasks in $\mathcal{T}_l$, employing effective labor $\exp\left(s_{i,\tau} + \varepsilon_{i,t}\right) \cdot \ell_{i,\tau,t}$ to produce task τ.[19] For any task $\tau \in \mathcal{T}_m$, the firm chooses what quantity of capital m_τ to rent. The firm thus optimizes output subject to the constraints

$$\sum_{\tau \in \mathcal{T}_l} \ell_{i,\tau,t} = 1$$
$$x_{i,\tau,t} = \exp\left(s_{i,\tau} + \varepsilon_{i,t}\right) \cdot \ell_{i,\tau,t} \quad \text{if } \tau \in \mathcal{T}_l$$
$$x_{i,\tau,t} = \exp\left(z_\tau\right) \cdot m_\tau \quad \text{if } \tau \in \mathcal{T}_m$$

Occupational choice. In every period t, each worker chooses an occupation to work in. Given their skill vector s_i and productivity shock $\varepsilon_{i,t}$, they fully anticipate their earnings conditional on entering occupation $o \in \mathcal{O}$. We assume that in any period t, the worker chooses the occupation yielding the highest utility given their individual vector of occupation-specific wages and preference shocks $u_{i,\cdot,t}$.[20] We further assume that each worker has log utility over their consumption of the numeraire, which is equal to their wage. Thus, the worker's occupational choice $\hat{o}_{i,t}$ is a function of log wages:

$$\hat{o}_{i,t} = \text{argmax}_o w_{i,o,t} + u_{i,o,t} \tag{2}$$

Discussion. We close this section by noting that the model is avowedly partial equilibrium in nature: We do not model households' demand for heterogeneous consumption goods and, instead, treat occupational output prices as fixed — equivalently, demand for occupational output is perfectly elastic. This represents a notable departure from a majority of the literature, as we abstract from a central force – i.e., the interplay between non-neutral technological progress and consumer preferences — that can dampen or amplify the wage effects of automation (Autor and Dorn, 2013). While we could relax this assumption, we deliberately impose it to isolate what

Section I). Our focus lies on the distributional effects.

[19] For ease of notation, we suppress human skills for machine tasks from the vector of human skills. Thus, $|\mathcal{T}_l| = n_{\text{skills}}$.

[20] We introduce no exogenous occupational switching frictions, so any persistence in occupational choices arises endogenously from the interaction of task-level skill specialization and occupational differences in task loadings.

we argue are important and heterogeneous labor market effects of automation that arise from induced shifts in task requirements within occupations, i.e., job transformation.

2.2 Optimality conditions and equilibrium

We next characterize optimality conditions, derive formulas for equilibrium wages and occupational choice, and then define an equilibrium.

Firm optimality and output. The firm's problem is

$$\max_{\ell_{i,o,\tau,t}, m_{i,o,\tau,t}} \prod_{\tau \in \mathcal{T}_l} \left(\exp(s_{i,\tau} + \varepsilon_{i,t})\ell_{i,o,\tau}\right)^{\alpha_{o,\tau}} \prod_{\tau \in \mathcal{T}_m} \left(\exp(z_\tau) m_{i,o,\tau,t}\right)^{\alpha_{o,\tau}} - \exp(w_{i,o,t}) \cdot 1 - r \sum_{\tau \in \mathcal{T}_m} m_{i,o,\tau,t}$$
$$\text{s.t.} \sum_{\tau \in \mathcal{T}_l} \ell_{i,o,\tau,t} = 1.$$

Defining $m_{i,o,t} := \sum_{\tau \in \mathcal{T}_m} m_{i,o,\tau,t}$, taking first order conditions yields

$$\ell_{i,o,\tau,t} = \frac{\alpha_{o,\tau}}{\sum_{\tau' \in \mathcal{T}_l} \alpha_{o,\tau'}} \forall \tau \in \mathcal{T}_l, \tag{3}$$

$$m_{i,o,\tau,t} = \frac{\alpha_{o,\tau}}{\sum_{\tau' \in \mathcal{T}_m} \alpha_{o,\tau'}} m_{i,o,t} \forall \tau \in \mathcal{T}_m, \tag{4}$$

$$\left(\sum_{\tau \in \mathcal{T}_m} \alpha_{o,\tau}\right) \frac{y_{i,o}}{m_{i,o,t}} = r. \tag{5}$$

which implies that log output equals

$$\log y_{i,o,t} = \left[\sum_{\tau \in \mathcal{T}} \frac{\alpha_{o,\tau}}{\sum_{\tau \in \mathcal{T}_l} \alpha_{o,\tau}} \log(\alpha_{o,\tau})\right] - \log\left(\sum_{\tau \in \mathcal{T}_l} \alpha_{o,\tau}\right) + \left[\sum_{\tau \in \mathcal{T}_m} \frac{\alpha_{o,\tau}}{\sum_{\tau \in \mathcal{T}_l} \alpha_{o,\tau}} (z_\tau - \log r)\right]$$
$$+ \left[\sum_{\tau \in \mathcal{T}_l} \frac{\alpha_{o,\tau}}{\sum_{\tau \in \mathcal{T}_l} \alpha_{o,\tau}} s_{i,\tau}\right] + \varepsilon_{i,t}.$$

Wages. Zero profits imply $w_{i,o,t} = \log\left(\sum_{\tau \in \mathcal{T}_l} \alpha_{o,\tau}\right) + \log y_{i,o,t}$, which means that the log wage of individual i in occupation o given their skill vector $s_{i,\cdot}$ and their productivity shock $\varepsilon_{i,t}$ can be written as the sum of an occupation-specific intercept, the weighted sum of log skills, and $\varepsilon_{i,t}$:

$$w_{i,o,t} = \mu_o + \sum_{\mathcal{T}_l} \frac{\alpha_{o,\tau}}{LS_o} \cdot s_{i,\tau} + \varepsilon_{i,t} \tag{6}$$

where $\mu_o = \sum_{\tau\in\mathcal{T}} \frac{\alpha_{o,\tau}}{\sum_{\tau\in\mathcal{T}_l}\alpha_{o,\tau}} \log\left(\alpha_{o,\tau}\right) + \left(\sum_{\tau\in\mathcal{T}_m} \frac{\alpha_{o,\tau}}{\sum_{\tau\in\mathcal{T}_l}\alpha_{o,\tau}} \left(z_\tau - \log r\right)\right)$ and $\text{LS}_o = \sum_{\tau\in\mathcal{T}_l}\alpha_{o,\tau}$ is the labor share in occupation o. For future reference, it is useful to define the following matrix of task weights.

Remark 1 (Task-weight matrix.)**.** *The matrix A, defined as*

$$A = \begin{pmatrix} \frac{\alpha_{1,1}}{LS_1} & \frac{\alpha_{1,2}}{LS_1} & \cdots & \frac{\alpha_{1,n_{skill}}}{LS_1} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{\alpha_{n_{occ},1}}{LS_{n_{occ}}} & \frac{\alpha_{n_{occ},2}}{LS_{n_{occ}}} & \cdots & \frac{\alpha_{n_{occ},n_{skill}}}{LS_{n_{occ}}} \end{pmatrix} \in \mathbb{R}^{n_{occ}\times n_{skill}} \tag{7}$$

summarizes the relative weights attached to each task $\tau \in \mathcal{T}_l$ across occupations $o \in \mathcal{O}$. The row vector $A_o := A_{o,\cdot}$ contains the task weights corresponding to occupation o.

Using this notation, we can write the vector of wages for a worker with skill vector s as

$$w = \mu + As + \varepsilon_{i,t} \in \mathbb{R}^{n_{\text{occ}}}$$

The wage equation allows for the following intuitive decomposition:

$$w_{i,o,t} = \mu_o + \underbrace{\frac{1}{n_{\text{skill}}}\sum_{\mathcal{T}_l} s_{i,\tau}}_{\text{scalar absolute advantage}} + \text{Cov}\left(n_{\text{skill}}\cdot\frac{\alpha_{o,\cdot}}{LS_o}, \underbrace{s_{i,\cdot} - \frac{1}{n_{\text{skill}}}\sum_{\mathcal{T}_l} s_{i,\tau}}_{\text{specialization vector}}\right) + \varepsilon_{i,t}$$

where the covariance operator is with respect to equal weights. The wage of a worker thus depends on their absolute advantage (captured by the average skill) and on how much the worker specializes in the skills that are important for that occupation (captured by the covariance term).

Occupational choice. Given the utility maximization problem in equation (2), the probability that individual i chooses occupation o conditional on their wage vector $w_{i,\cdot}$ is

$$P(\hat{o} = o|w_{i,\cdot}) = \frac{\exp(w_{i,o}/\nu)}{\sum_{o'}\exp(w_{i,o'}/\nu)} \tag{8}$$

Equilibrium. An equilibrium is defined as a joint distribution Γ of occupation choices, log wages w, log skills s and idiosyncratic productivity shocks $\varepsilon_\cdot$, such that: (i) equation (6) holds at any point in the distribution (that is, firms make zero profits); (ii) the marginal distribution of occupations conditional on wages follows equation (8) (that is, workers optimize); and (iii) the unconditional marginal distributions of skills s and occupational shocks ε follow $\mathcal{N}(\bar{s}, \Sigma_s)$ and

$\mathcal{N}(0, \varsigma^2 I)$, respectively.

2.3 The wage effects of automation in theory

We now describe how automation of a particular task $\tau^\star$ is formalized. To do this, we endogenize the task assignment $(\mathcal{T}_l, \mathcal{T}_m)$, then characterize the induced wage change as a function of skills. We allow for arbitrarily large shocks with potentially non-linear effects rather than relying on first-order perturbation methods, which may not capture a shock's transformative nature.

2.3.1 Automation in the model

To conceptualize automation and its wage effects we now endogenize the task assignment $(\mathcal{T}_l, \mathcal{T}_m)$ and make it dependent on the underlying machine productivity z_τ at every task τ. Appendix A.1 outlines the set of assumptions we introduce to this end and derives a minimum-machine productivity threshold $\bar{z}_{\tau^\star}$ above which automation optimally occurs in equilibrium.[21]

We model an automation shock as a one-time, permanent change of $z_{\tau^\star}$ that leads to the reassignment of task $\tau^\star$ from labor to machines. Let the prime symbol denote variables after an automation shock in period $t^\star$ and let $(\mathcal{T}_l, \mathcal{T}_m)$ be the initial task allocation. The new task sets are thus

$$\mathcal{T}_l' = \mathcal{T}_l \setminus \{\tau^\star\}, \quad \mathcal{T}_m' = \mathcal{T}_m \cup \{\tau^\star\}.$$

The qualitative analyses in this Section 2.3 hold for any value of machine productivity $\{z_{\tau^\star} : z_{\tau^\star} \geq \bar{z}_{\tau^\star}\}$. For our quantitative analyses in Section 4 we will need to take a stand on the exact value of $z_{\tau^\star}$; we will consider the specific case where machine productivity at the newly automated task is *just* high enough to make the automation of task $\tau^\star$ optimal from the perspective of all firms, i.e., $z_{\tau^\star} = \bar{z}_{\tau^\star}$.

Associated with the automation shock is a change in the occupational task weight matrix A, as defined in Remark 1, whereby automation reduces the weight on the automated task to zero

[21] In models of automation which do not explicitly feature task bundling or occupational choice, such as (Acemoglu and Restrepo, 2018b) and Acemoglu and Restrepo (2022), tasks can be ordered by the relative productivity of humans to machines. The threshold at which automation occurs can then be written as the point at which this productivity ratio equals the ratio of wages to capital costs. The introduction of occupational choice and skill heterogeneity in our setting complicates this simple rule. To maintain tractability, we maintain the assumption that the automation choice is common across occupations and workers of different skills.

and increases the weight on all other entries proportional to their weight:

$$\begin{aligned} A'_o - A_o &= \left(\frac{\alpha_{o,1}}{LS'_o} \cdot \frac{\alpha_{o,\tau^\star}}{LS_o} \quad \frac{\alpha_{o,2}}{LS'_o} \cdot \frac{\alpha_{o,\tau^\star}}{LS_o} \quad \dots \quad -\frac{\alpha_{o,\tau^\star}}{LS_o} \quad \dots \right) \\ &= \left(\frac{\alpha_{o,1}}{LS'_o} \quad \frac{\alpha_{o,2}}{LS'_o} \quad \dots \quad -1 \quad \dots \right) \cdot \frac{\alpha_{o,\tau^\star}}{LS_o} \\ &= (A'_o - \iota_{\tau^\star}) \cdot \frac{\alpha_{o,\tau^\star}}{LS_o} \end{aligned}$$

The change in the occupational wage vector of a worker i arises is a function of this change:

$$\begin{aligned} \Delta w_{i,o} &= w'_{i,o} - w_{i,o} = \Delta\mu_o + (A'_o - A_o)s_i + \Delta\varepsilon_i \\ &= \Delta\mu_o + \frac{\alpha_{o,\tau^\star}}{LS_o} \left(\sum_{\mathcal{T}_l \setminus \tau^\star} \frac{\alpha_{o,\tau}}{LS_o - \alpha_{o,\tau^\star}} s_{i,\tau} - s_{i,\tau^\star} \right) + \Delta\varepsilon_i \end{aligned} \tag{9}$$

where

$$\Delta\mu_o = \frac{\alpha_{o,\tau^\star}}{LS_o - \alpha_{o,\tau^\star}} \left(z_{\tau^\star} - \log r + \mu_o \right).$$

Equation (9) captures two important terms. First, workers are more likely to see increases in their origin-occupation wage when $\Delta\mu_o$ is large. This effect captures both the negative displacement and the positive productivity effects which together fully characterize the effects of automation in the canonical task-based model. In our model, a second important effect shapes earnings responses: the shift in task weights, which we refer to as *job transformation*. This second component depends on the worker's task-specific skills: Workers are more likely to benefit if they are relatively unskilled in automated task relative to other tasks, where the latter are weighted by the loadings of their current occupation.[22]

Partial automation. This way of conceptualizing automation nests the case where a skill becomes *partially* automated. A skill is said to be partially automated when A_o changes in a way that does not set the loading of the automated skill to zero:

$$A'_o - A_o = (A'_o - \iota_{\tau^\star}) \cdot \frac{\alpha_{o,\tau^\star}}{LS_o} \cdot \delta_{\tau^\star}, \quad \delta_{\tau^\star} \in (0, 1)$$

To interpret this case, suppose $\tau^\star$ comprises two distinct tasks instead of one. We effectively suppose workers' skills for these two tasks are always identical – that is, perfectly correlated with

[22] This mechanism mirrors the analysis of team complementarities under skill specialization in Freund (2023): Your productivity is enhanced by a coworker – whether human or artificial – insofar as their presence enables you to focus on those task you are best at; and the magnitude of this complementarity effect is increasing in the degree of skill specialization.

identical skill means and standard deviations – but only one of the two tasks can be automated.

2.3.2 The role of task bundling

We just saw that in our model automation affects wages by altering the task content of occupations. We next explain that these effects distinctively arise when the economy exhibits *task bundling*; that is, multiple tasks are performed concurrently within the same occupation. We use the following nomenclature:

Remark 2 (Task bundling.)**.** *An occupation features* ***task-bundling*** *if*

$$|\{\tau \in \mathcal{T}_l : \alpha_{o,\tau} > 0\}| > 1.$$

Conversely, the economy features a ***no-bundling property*** *if no occupation features task bundling:*

$$|\{\tau \in \mathcal{T}_l : \alpha_{o,\tau} > 0\}| = 1 \quad \forall o \in \mathcal{O}.$$

In a no-bundling economy, there exists an assignment function $g : \mathcal{O} \to \mathcal{T}$ that pins down the unique task required in any given occupation.[23] In this case, the wage equation reduces to

$$w_{i,o,t} = \mu_o + s_{i,g(o)} + \epsilon_{i,o,t}. \tag{10}$$

In a no-bundling economy, workers in an occupation o subject to automation thus experience wage changes that are solely driven by changes in the occupation-specific shifter, i.e. $\Delta\mu_o$. The wage changes are, thus, driven by the well-understood balance between negative displacement effect, associated with a declining labor share, and positive productivity effects, driven by $\bar{z}_{\tau\star}$.[24] Crucially, workers do not experience any effects from a changing task mix of their occupation. Moreover, conditional on staying in the same occupation, all workers in an occupation experience the same wage change.

In contrast, under task bundling, individual wages change also because automation shifts the task content of their occupation. Equation (9) describes these effects for an individual worker. The next section offers a characterization of these wage effects due to job-transformation.

[23] A special case of this is the case where $A = I$. With this case, our model nests the standard Roy model. Note that occupations having the same A_o is not sufficient for them to be perfect substitutes from a worker's perspective, as they may involve labor shares or different machine tasks with differing productivities.

[24] For a detailed review see Acemoglu *et al.* (2025). A subtle difference in the operation of the positive productivity effects compared to the canonical model is worth noting. For example, in Acemoglu and Restrepo (2022, cf. equations (6) and (13)), the productivity effect raises the wages of *all* workers. What underlies this feature is the assumption that substitution across all tasks is governed by a uniform elasticity parameter. In contrast, in our model, production occurs at the level of occupations, so automation carries no positive productivity effects for occupations that do not utilize the automated task.

2.3.3 The distributional effects of automation

We next characterize the economy-wide distributional effects of automation in the presence of task bundling by deriving a transparent decomposition of occupation-level wage changes due to automation.[25]

Remark 3 (Occupational wage change.)**.** *The occupation-level average wage change arising from automation is*

$$\begin{aligned}
&\mathbb{E}[w'_o|\hat{o}'=o]-\mathbb{E}[w_o|\hat{o}=o]\\
&=\overbrace{\mathbb{E}[w'_o|\hat{o}=o]-\mathbb{E}[w_o|\hat{o}=o]}^{\Delta w_o \text{ of incumbents}}+\overbrace{\mathbb{E}[w'_o|\hat{o}'=o]-\mathbb{E}[w'_o|\hat{o}=o]}^{\text{re-sorting}}\\
&=\overbrace{\underbrace{\Delta\mu_o}_{\text{productivity and displacement}}+\underbrace{(A'_o-A_o)\cdot\bar{s}}_{\text{task shift}}+\underbrace{(A'_o-A_o)(\bar{s}_{|o}-\bar{s})}_{\text{selection}}}^{\Delta w_o \text{ of incumbents}}\\
&\quad+\underbrace{\mathbb{E}[w'_o|\hat{o}'=o]-\mathbb{E}[w'_o|\hat{o}=o]}_{\text{re-sorting}}
\end{aligned}$$

and thus

$$\begin{aligned}
&\mathbb{E}[w'_o|\hat{o}'=o]-\mathbb{E}[w_o|\hat{o}=o]\\
&=\overbrace{\underbrace{\Delta\mu_o}_{\text{productivity and displacement}}+\underbrace{(A'_o-A_o)\cdot\bar{s}}_{\text{task shift}}+\underbrace{\nu^{-1}(A'_o-A_o)\Sigma\left(A_o^\intercal-\sum_{o''}h_{o''}(\bar{s}_{|o})A_{o''}^\intercal\right)}_{\text{selection}}}^{\Delta w_o \text{ of incumbents}}\\
&\quad+\underbrace{\nu^{-1}A'_o\Sigma\left(\left((A'_o-A_o)^\intercal-\sum_{o''}\left(h'_{o''}(\bar{s}'_{|o})(A'_{o''})^\intercal-h_{o''}(\bar{s}_{|o})A_{o''}^\intercal\right)\right)\right)}_{\text{re-sorting}}. \qquad (11)
\end{aligned}$$

[25] We rely on Laplace approximations to derive our results in this section. This method approximates a posterior distribution with an appropriately chosen multivariate normal distribution. Details are in Appendix A.2.

where

$$\bar{s}_{|o} = \bar{s} + \nu^{-1}\Sigma \overbrace{\left(A_o^\intercal - \sum_{o''} h_{o''}(\bar{s}_{|o}) A_{o''}^\intercal\right)}^{\text{relative task intensity of occupation } o} \tag{12}$$

is the posterior skill mean of workers choosing occupation o and

$$h_o(s) = \frac{\exp(\nu^{-1}\mu_{o'} + \nu^{-1}A_{o'} \cdot s)}{\sum_{o''} \exp(\nu^{-1}\mu_{o''} + \nu^{-1}A_{o''} \cdot s)} \tag{13}$$

is the employment share of occupation o for workers with skill s.

The first line of equation (11) captures occupational wage effects when worker composition is held constant. The second line captures wage changes from worker re-sorting into and out of occupation o. We discuss each term individually.

- The first term captures the net impact of the standard *displacement and productivity effects*. These effects operate even without task bundling, whereas the remaining terms are unique to economies with task bundling.

- The *task shift* effect captures how shifts in occupational task weights alter the productivity of an average worker and therefore wages. These effects can be positive when the task composition of the occupation shifts to more productive tasks ("task upgrading") or negative otherwise ("task downgrading").

- The *selection* term captures that incumbent workers in occupation o may have skills different from the population average and therefore respond differently to automation. Equation (12) characterizes the posterior mean skill of those choosing occupation o: it exceeds $\bar{s}$ for those tasks used more intensively in o, especially for tasks with high skill variance. The selection effect tends to be negative for more exposed occupations, since for these occupations A_o has larger positive entries exactly where $(A'_o - A_o)$ is negative. This effect strengthens when selection into automated skills is stronger, which occurs when ν is low and when skills at the automated task are more dispersed across workers.

- The *re-sorting* effect captures that workers choosing occupation o after automation may have a different skill distribution than before. For exposed occupations, the sign is generally positive under full automation. The effect strengthens when skills, especially heavily utilized non-automated skills, are more dispersed. It can also be more pronounced for

more dispersed automated skills when they correlate negatively with other skills heavily utilized in affected occupations.

This decomposition illustrates that the effects of automation, especially as they relate to the implications of task bundling, depend in a critical way on the underlying distribution of worker skills s; that is, both the skill means $\bar{s}$ and the co-variance matrix Σ_s are informative about the properties of different types of automation shocks. A quantitative assessment of such shocks thus requires knowing the empirical values of these objects. The challenge is that the multi-dimensional skill distribution is generally unobserved. However, as we show in the following section, our model provides a natural framework to estimate this skill distribution.

3 Theory Meets Data

We now take the theoretical model to data. We begin by describing the theoretical foundations for our approach to estimating the model's parameters (Section 3.1), then explain its empirical implementation (Section 3.2), validate it in Monte-Carlo exercises (Section 3.3), and finally present the estimation results alongside a comparison of the model along a few important empirical dimensions (Section 3.4).

3.1 Estimation methodology

To estimate the parameters of the model, the following data are required: (i) a panel of workers, indexed by i, for whom both occupational choices $\hat{o}_{i,t}$ and wages $w_{i,\hat{o}_{i,t},t}$ are observed over time; (ii) the occupational task weight matrix A, as defined in Remark 1; as well as (iii) a measure of occupation-level labor shares LS_o. While (i) and (iii) are relatively weak data requirements[26], (ii) is more involved. We argue that it is possible to obtain good measures of (iii) directly from the data and discuss the construction of our empirical measure of A in Section 3.2.

Conditional on observing (i)-(iii), we make two identifying assumptions. First, we assume that the model is in steady state throughout our estimation window. Second, in the initial steady state there is only one composite machine task with productivity normalized to $\log r$.[27] This

[26] We discuss the construction of occupation-level labor shares in Appendix B.4.

[27] This assumption is necessary to identify skill means $\bar{s}$ in the population. In a model with an arbitrary number of unobserved machine tasks with unobserved productivity, $\bar{s}$ is not separately identified: Intuitively, occupations could offer high wages because they load on tasks with higher average skill or because they load on machine tasks with higher average productivity.

implies that the intercept term

$$\begin{aligned}\mu_o &= \sum_{\tau\in\mathcal{T}} \frac{\alpha_{o,\tau}}{\sum_{\tau\in\mathcal{T}_l}\alpha_{o,\tau}} \log\left(\alpha_{o,\tau}\right) + \left(\sum_{\tau\in\mathcal{T}_m} \frac{\alpha_{o,\tau}}{\sum_{\tau\in\mathcal{T}_l}\alpha_{o,\tau}} \left(z_\tau - \log r\right)\right) \\ &= \sum_{\tau\in\mathcal{T}_l} A_{o,\tau} \log\left(A_{o,\tau}\cdot LS_o\right) + \frac{1-LS_o}{LS_o}\left(\log LS_o\right)\end{aligned}$$

depends only on occupational labor shares $LS_o = \sum_{\tau\in\mathcal{T}_l}\alpha_{o,\tau}$ and elements of A and is therefore observable for all occupations.

Under these conditions, we can estimate the model with maximum likelihood techniques. In what follows, let $\hat{o}_{i,t}$ denote the recorded occupation choice of worker i in period t and $-\hat{o}_{i,t}$ be the set of occupations not chosen in period t. Let A be the matrix defined in (7). For a given worker observed in occupations $(\hat{o}_{o,1},\ldots,\hat{o}_{i,T})$ and earning wages $(w_{i,\hat{o}_1,1},\ldots,w_{i,\hat{o}_T,T})$,

$$\begin{pmatrix} s_1 \\ \vdots \\ s_{n_{\text{skill}}} \\ w_{i,\hat{o}_1,1} \\ \vdots \\ w_{i,\hat{o}_T,T} \end{pmatrix} = \begin{pmatrix} 0 \\ \vdots \\ 0 \\ \mu_{\hat{o}_1} \\ \vdots \\ \mu_{\hat{o}_T} \end{pmatrix} + \begin{pmatrix} I & 0 \\ A_{(\hat{o},\ldots,\hat{o}_T),\cdot} & I \end{pmatrix} \cdot \begin{pmatrix} s_1 \\ \vdots \\ s_{n_{\text{skill}}} \\ \varepsilon_{i,1} \\ \vdots \\ \varepsilon_{i,T} \end{pmatrix}, \quad \text{where} \quad \begin{pmatrix} s_1 \\ \vdots \\ s_{n_{\text{skill}}} \\ \varepsilon_{i,1} \\ \vdots \\ \varepsilon_{i,T} \end{pmatrix} \sim \mathcal{N}\left(\begin{pmatrix} \bar{s}_1 \\ \vdots \\ \bar{s}_{n_{\text{skill}}} \\ 0 \\ \vdots \\ 0 \end{pmatrix}, \begin{pmatrix} \Sigma_s & 0 \\ 0 & \varsigma^2 I \end{pmatrix}\right).$$

Thus, $w_{i,\hat{o}_{i,\cdot},\cdot}$ and s_i are jointly normal and thus yields easy to compute formulas for the distribution of $s_i|w_{i,\hat{o}_{i,\cdot},\cdot}$. The likelihood of observing $(w_{i,\hat{o}_{i,\cdot},\cdot},\hat{o}_{i,\cdot})$ is then given by

$$\begin{aligned}\mathcal{L}(w_{i,\hat{o}_{i,\cdot},\cdot},\hat{o}_{i,\cdot}|\nu,\varsigma,\bar{s},\Sigma_s) = \prod_i \int_s \Bigg[\left(\int_{w_{i,-\hat{o}_{\cdot,\cdot}}} \prod_t P(\hat{o}_{i,t}|w_{i,\cdot,t},\nu)\cdot f(w_{i,-\hat{o}_{i,t},t}|s,w_{i,\hat{o}_{i,\cdot},\cdot},\varsigma)\right) \\ \cdot f(s|w_{i,\hat{o}_{i,\cdot},\cdot},\varsigma,\bar{s},\Sigma_s)\Bigg]\cdot f(w_{i,\hat{o}_{i,\cdot},\cdot}|\varsigma,\bar{s},\Sigma_s)\end{aligned}$$

This expression involves a high-dimensional integral which makes it intractable. To overcome this challenge, we use Monte Carlo integration to compute a numerical approximation of the likelihood instead of evaluating this expression analytically. That is, instead of maximizing the analytical likelihood, we instead maximize the mean of a simulated statistical object that converges to the likelihood value for large sample sizes, n_0. It can be shown that, as $n_0 \to \infty$, the argmax of this object converges to the true maximum likelihood estimate under mild regularity conditions. We find that, in practice, $n_0 = 40$ yields a sufficiently accurate approximation to

deliver satisfactory results in a Monte Carlo exercise, which we report below.

Concretely, our implementation of this idea is as follows: For all individual workers i, we generate n_0 draws from

$$f(w_{i,-\hat{o}_{i,\cdot},\cdot}|w_{i,\hat{o}_{i,\cdot},\cdot},\varsigma,\bar{s},\Sigma_s)=\int_s f(w_{i,-\hat{o}_{i,\cdot},\cdot}|s,w_{i,\cdot,\hat{o}_{i,\cdot}},\varsigma)f(s|w_{i,\hat{o}_{i,\cdot},\cdot},\varsigma,\bar{s},\Sigma_s).$$

These draws can be generated by (i) drawing from the distribution $s_i|w_{i,\hat{o}_{i,\cdot},\cdot}$, (ii) computing $\varepsilon_{i,t}$ for every period (as a deterministic function of s_i and $w_{i,\hat{o}_{i,t},t}$), and (iii) computing the resulting vector of all occupational wages in every period. Using these wages, we then evaluate the mean of $P(\hat{o}_{i,t}|w_{i,\hat{o}_{i,\cdot},\cdot},\nu)$ to obtain an estimator for $\mathcal{L}_i(\theta)$:

$$\hat{\mathcal{L}}_i(w_{i,\hat{o}_{i,\cdot},\cdot},\hat{o}_{i,\cdot}|\nu,\varsigma,\bar{s},\Sigma_s)=\left(\frac{1}{n_0}\sum_j\prod_t P(\hat{o}_{i,t}|w_{j,\cdot,t},\nu)\right)\cdot f(w_{i,\hat{o}_{i,\cdot},\cdot}|\varsigma,\bar{s},\Sigma_s)$$

Holding constant all random variables in the estimator, we then proceed to maximize this objects over the parameter space $\theta=(\bar{s},\Sigma_s,\nu,\varrho)$. This parameter space is large and requires efficient numerical optimization methods. We utilize stochastic gradient descent paired with auto-differentiation techniques that allow us to efficiently compute gradients of $\hat{\mathcal{L}}_i$. Details are relegated to Appendix C.1.

A few clarifying comments on this identification approach are in order. First, it may be fair to ask whether the model parameters are even fully identified from observable data. Appendix A.3 shows that the answer to this question is yes and provides an identification proof. Second, it is useful to consider what features of the model ultimately lead to this identification result and make our identification strategy viable. The answer to this lies in the assumed connection between observable data (occupational choices and wages) and the estimation object (the distribution of skills): Both observed and unobserved wages link to skills through the linear wage equation and the time-share matrix A, which, importantly, we assume to be known. Occupational choices in every period further place restrictions on unobserved wages through the observed occupational choices of a worker. Note that identification does not rely on the order in which occupational choices are observed (the sequence of occupations does not enter the likelihood) but instead only relies on the total number of times an occupation is observed for a given worker. That said, to the extent that occupational choices may be subject to switching frictions in reality, for time-limited samples we are likely to somewhat overstate skill specificity in our baseline model, which does not feature such frictions.

3.2 Data & measurement

As noted, we require three data sources: a worker panel with information on wages and occupational choices, the occupational task weight matrix A, and occupational labor shares. For the worker panel we use the National Longitudinal Survey of Youth (NLSY) 1979. We construct the task weight matrix from a list of detailed occupational tasks from O*NET and by leveraging natural language processing (NLP) tools as well as large language models (LLMs). Lastly, we construct occupational labor shares using data from the Bureau of Economic Analysis (BEA). This section explains data sources and processing in detail.

NLSY. The NLSY 1979 tracks 6,033 workers' occupations and wages between 1979 and 2018, comprising 110,618 total observations. We construct an annual panel comprising each individual's primary job (if any), where in the case of multiple jobs the "primary job" is selected based on weekly hours worked. Wages are deflated using the CPI (1982-1984=100). We drop individuals in the military sample and the minority oversample. Following Lise and Postel-Vinay (2020), we create a harmonized occupational classification at the SOC-2000 level using crosswalks from Sanders (2012). We use the "minor groups" of occupations, of which there are 93 in our data.

Tasks & occupational task weights. We construct the occupational task weight matrix A defined in Remark 1 in two main steps. In step one, we cluster approximately 20,000 detailed tasks with similar skill requirements using NLP. In step two, we measure the weights different occupations attach to these task clusters using LLMs. Figure 1 summarizes this workflow, and Figure 2 illustrates the mapping from tasks to clusters with examples. Appendices B.2 and B.3 contain further details on the use of clustering algorithms and LLMs.

Our starting point is the list of 18,796 detailed, occupation-specific tasks in the O*NET database (version 29.2). Using this granular list of tasks is instrumental for our purposes, as many technology-specific task automation exposure measures use this list as a reference point, enabling us to use these exposure measures in Section 4.1 to pin down technology- and task-specific automation shocks. In step one, we group these detailed tasks into clusters that serve as the empirical analogue to the set of human-performed tasks $\mathcal{T}_l$. Our methodology aims to group together tasks with similar skill requirements — if a person is proficient in one detailed task from cluster τ, they should also be able to perform another task assigned to τ equally well. The constraint we confront is that O*NET provides limited meta-data, and specifically no skill requirements, for the detailed tasks.[28]

[28]We explored several alternatives, including clustering on the raw embeddings for the detailed task statements. In practice, however, this approach leads to clusters of tasks that share a similar *context* (e.g., "hospital") but have very different *skill requirements*. For example:"Provide and manage long-term, comprehensive medical care, including diagnosis and nonsurgical treatment of diseases, for adult patients in an office or hospital." versus "Report facts concerning accidents or emergencies to hospital personnel or law enforcement officials."

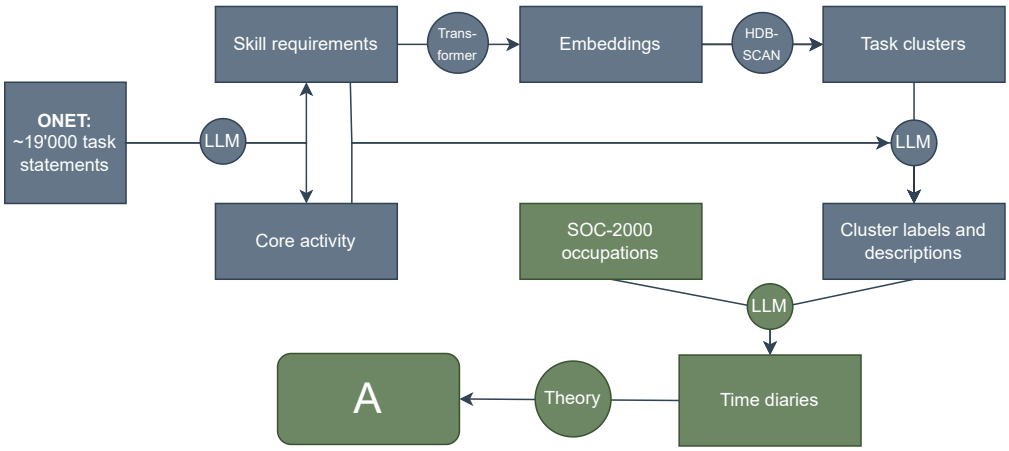


Figure 1: Schematic overview of the measurement of tasks and the A matrix

Notes. Step 1 (colored in blue) involves the clustering of tasks, step 2 (colored in green) the measurement of occupational task weights.

We process these task statements using a multi-stage NLP pipeline. We start by feeding them through OpenAI's ChatGPT-4o model to identify the core activity and the 1-5 essential skills required to be productive in the task. Next, we create word embeddings of these skill requirement descriptions —representations of their semantic content in a high-dimensional vector space — using Alibaba's GTE transformer model (gte-Qwen2-1.5B-instruct). This enables us to cluster detailed tasks by applying a hierarchical density-based clustering algorithm (HDBSCAN, McInnes *et al.* (2017)) to the embeddings. This yields a set of 38 clusters.[29]

To interpret the resulting clusters, we use OpenAI's o3-mini-high model, feeding it representative examples from each cluster and prompting it to generate descriptive labels and concise summaries that capture the underlying skill requirements. Appendix Table B.2 gives examples of the core activity and skill requirements extracted for several detailed tasks as well as the resulting cluster. Appendix Table B.3 lists all 38 task clusters.

In step two, we then construct the occupational task weight matrix A. The theory guides measurement, as equation (3) indicates that each entry of the A matrix corresponds to the optimally chosen share of time allocated to task τ in any occupation o. We measure these shares across our 38 task categories – which we describe in terms of labels as well as the summary descriptions derived above – by prompting OpenAI's o3-mini-high model to construct time allocation diaries for each occupation. Appendix B.3 details our LLM prompts, while Appendix

[29]Two remarks about the HDBSCAN algorithm are in order. First, tasks that don't contribute to any stable cluster are automatically classified as noise; we drop these tasks from our analysis. Second, the algorithm automatically determines the number of clusters k through a hierarchical approach based on cluster stability, unlike the familiar k-means algorithm where k is a user input.

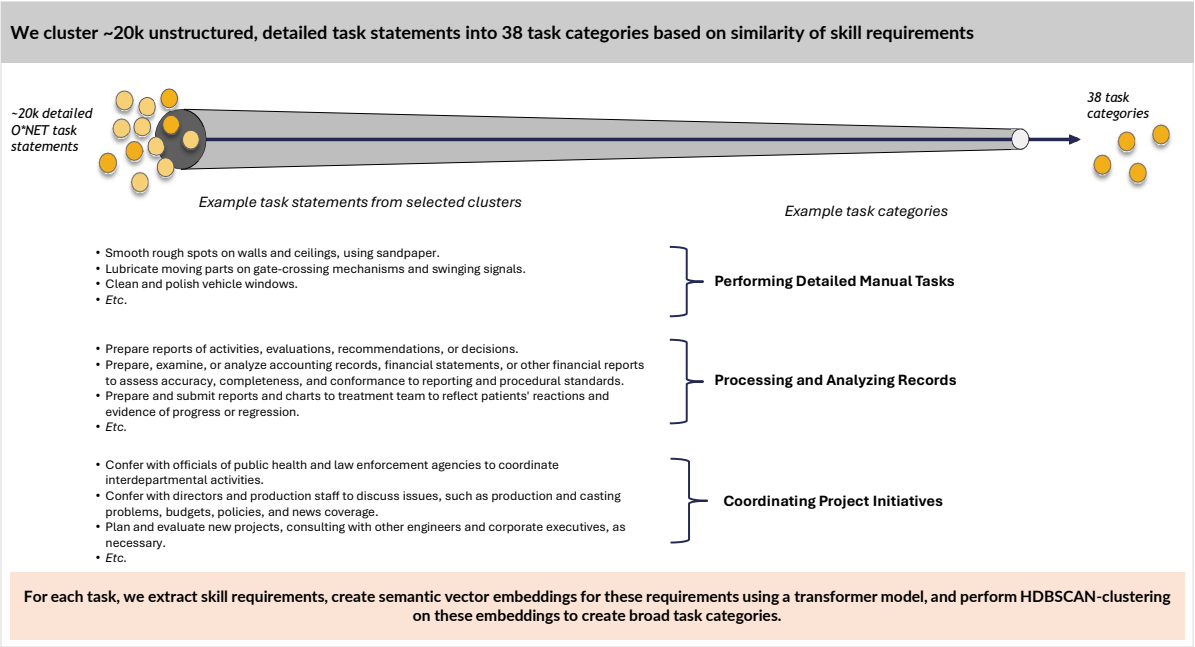


Figure 2: Examples of mapping from detailed tasks to clusters

C.2.4 discusses an alternative way of constructing A from O*NET task importance weights.

The A matrix thus obtained corroborates the importance of task bundling and has intuitive properties. Only two occupations have a single task comprising more than half of incumbents' time, and in fewer than 30% of occupations does a single task account for more than a quarter of total time. Figure 3 visualizes A. For visualization purposes, we bi-cluster tasks and occupations, grouping similar categories together, as indicated by the white dividing lines. The A matrix exhibits sparsity with intuitively aligned entries. For instance, "Performing detailed manual tasks" appears as a prominent task across both service sector occupations, like "Food and Beverage Serving Workers" and manufacturing roles, like "Assemblers and fabricators." However, the former, service-sector occupations additionally emphasize tasks like "Providing customer service" while the latter, manufacturing-oriented jobs involve more technical tasks like "Operating, Calibrating, and Inspecting Equipment."

This LLM-powered approach is a useful and flexible tool, but of course invites some immediate questions: How could the LLM know this information? And are the resulting measurements reliable? Regarding the first question, LLM training data comprises virtually the entire internet, including vast amounts of unstructured data on what people across different occupations do

at work, as well as summaries of time diary surveys reported in research papers. Since these data sources generally do not reference our exact tasks or occupations, and much input data is qualitative, the LLM's quantitative output results from interpolation.[30] Given the black-box nature of this data construction, we conduct a battery of exercises to evaluate LLM capability in constructing occupational time diaries. We summarize them here and refer to appendix B.3.2 for further information.

First, we compare LLM-generated task weights at the occupation-cluster level to the average importance rating that O*NET assigns to detailed tasks within each cluster. While O*NET importance weights do not directly map onto our A matrix entries – whereas time shares do, per equation (3) – they are strongly correlated with our baseline measures. Second, we exploit a unique 2012 supplemental survey by the German Federal Institute for Vocational Training (Bundesinstitut fuer Berufsbildung, BIBB) in which workers across many occupations report their time allocation across 17 tasks. Though the occupations and tasks differ from our baseline analysis, our LLM-based method is flexible enough to generate time diaries for German BIBB classifications. This comparison reveals highly correlated time shares between the two approaches. Third, we use O*NET's Generalized Work Activities as a task classification, with importance ratings as weights. LLM-generated time shares for these activities again align strongly with importance ratings. Finally, we establish LLM internal consistency: occupational task weights constructed by averaging across constituent minor categories are highly correlated with those derived by directly querying the model for major groups. Together, these validation exercises confirm that our LLM-based approach aligns with established measurement frameworks while offering greater flexibility, notably in task and occupation classifications.

Occupation-level labor shares. Finally, we construct LS_o by combining industry-level data on value-added and wage payments from the BEA-BLS Integrated Industry-level Production Accounts with information on wage payments at the industry-occupation level from the BLS Occupational Employment and Wage Statistics (OEWS) Tables. Concretely, as value-added and hence labor shares are generally defined and measured at the establishment- or industry-level, we apportion industry-level value-added to occupations based on their share in an industry's total wage bill. Appendix B.4 provides more details.

3.3 Validation: Monte-Carlo exercises

We are now ready to estimate parameters according to the approach outlined in Section 3.1. To show that this approach indeed robustly identifies the parameters, we preface our estimation

[30] Our LLM usage resembles employing a vast pool of research assistants with unlimited time to collect diverse data sources and make judgment calls in translating them into cardinal time shares.

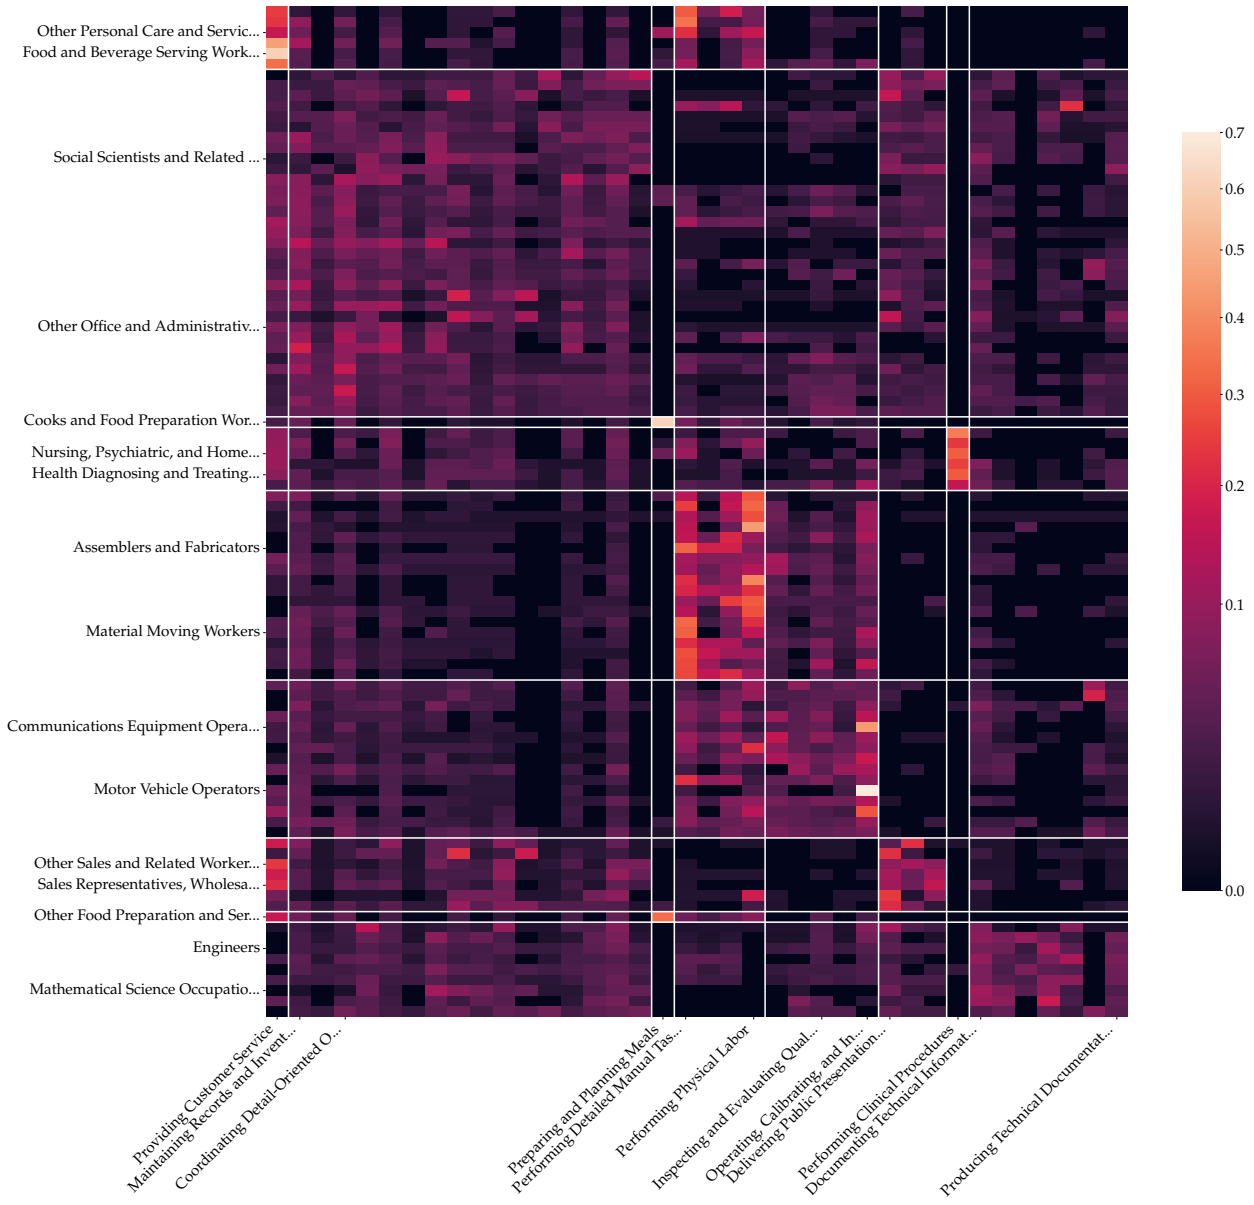


Figure 3: Task weight matrix

Notes. This figure shows the measured A matrix; each cell value corresponds to $\frac{\alpha_{o,\tau}}{\sum_{\tau \in \mathcal{T}_l} \alpha_{o,\tau}}$. To aid visualization, the matrix is reordered using a spectral co-clustering algorithm, and example tasks and occupations are highlighted.

results by conducting a Monte-Carlo exercise. This exercise involves three steps: First, we generate parameter estimates by applying the methodology described above to the real world data described in Section 3.2. Second, we generate artificial data from our estimated model under the parameters estimated in step one. Third, we apply our methodology once more to the artificially generated data, then compare the resulting estimates with the estimates generated in step one. If our method correctly recovers the data generating parameters, then the parameters estimated in step 3 should align well with the parameters estimated in step 1, which are used as inputs when generating the artificial data.

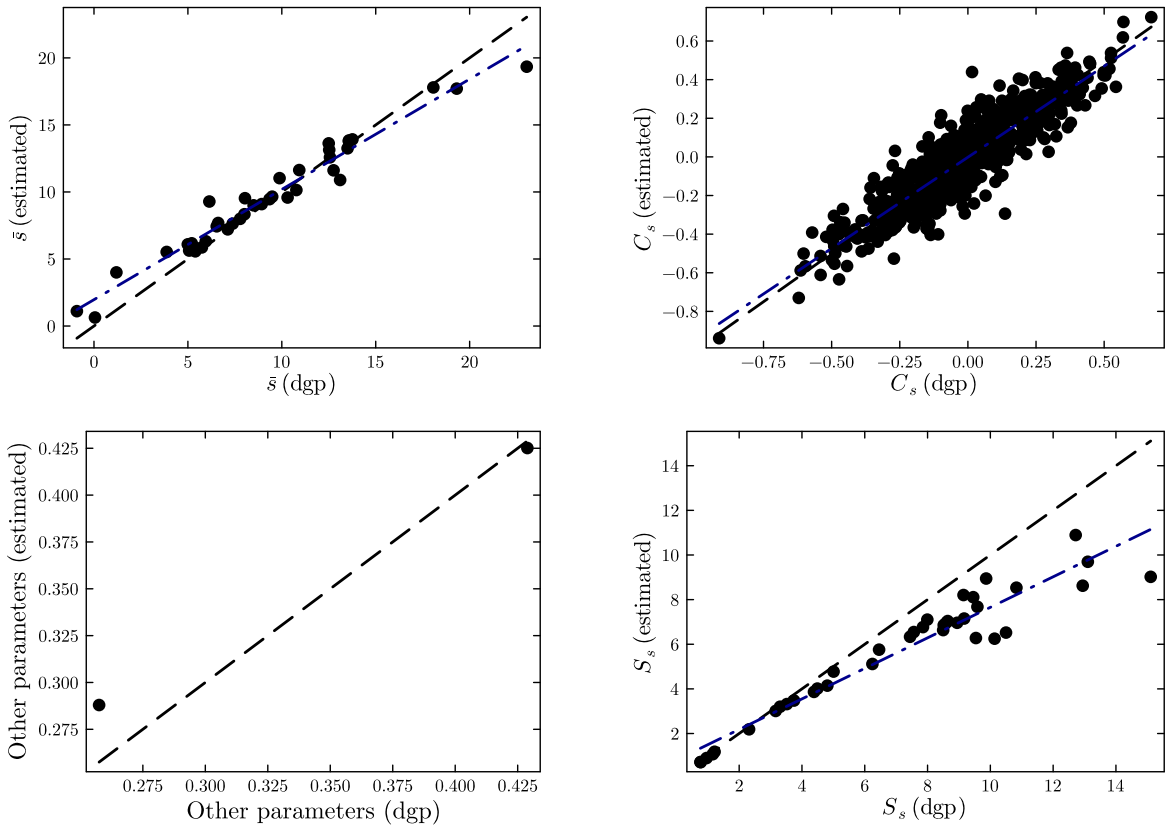


Figure 4: Data generating parameters and their estimates

Notes: The horizontal axis displays parameter values used to generate artificial data. The vertical axis displays corresponding estimated values. "Other parameters" refers to the tuple (ν, ς). The black dashed line is the 45 degree line. The blue dash-dotted line is the line of best fit.

This exercise corroborates our methodology, with the parameters estimated aligning well with the data generating process in the simulation ("dgp"). Figure 4 illustrates this comparison, with each panel showing one set of estimated parameters. We split the skill covariance matrix into its correlation component (C_s, omitting the diagonal of ones in the figure) and its vector of standard deviations S_s; that is, we decompose it according to $\Sigma_s = \text{diagm}(S_s) \cdot C_s \cdot \text{diagm}(S_s)$.

The remaining two panels show the estimated and data generating parameters for the vector of mean log skills, $\bar{s}$, and the remaining parameters (ν, ς), respectively. The fit is generally good; in particular, we are able to capture the rather large number of parameters governing bi-lateral skill-correlations quite well.

3.4 Estimation results & model validation

3.4.1 Parameter estimates

For the scalar parameters, we estimate $\nu = 0.26$ and $\varsigma = 0.43$. The estimate of ν implies that reducing prospective wages in a given occupation by 1% lowers the odds of choosing this occupation by about 3.8% since $\frac{1}{\nu} \approx 3.8$. $\varsigma = 0.43$ indicates that a one-standard-deviation occupation-specific random productivity shock can raise or lower wages by about 43% in a given year.[31]

Next, we investigate the mean and dispersion of skills. Figure 5 documents the estimate of mean skills $\bar{s}$ and illustrates their dispersion as error bands, which is given by their standard deviation (S_s where $\Sigma_s = \text{diagm}(S_s) \cdot C_s \cdot \text{diagm}(S_s)$). The values of $\bar{s}$ indicate tasks that are more or less productive on average in the population. For example, “coordinating project initiatives" and “monitoring and inspecting systems" are tasks that are more productive on average, whereas for example, “mediating and consulting clients" is associated with lower average skill. However, skills also differ in their dispersion across the worker population. For example, workers differ most in their capability to “analyze and optimize systems” and in their ability to “produce technical documentation”, but relatively little in their skill for “performing detailed manual tasks”, “preparing and planning meals”, or “performing physical labor”.

This contrast is interesting from the perspective of the formulas derived in Section 2.3.3. On one hand, the decomposition of equation (11) indicates that the magnitude of some automation effects, such as selection and re-sorting effects, depend crucially on the between-worker skill dispersion at the affected task. Figure 5, on the other hand, indicates that the dispersion of skills varies widely with the type of task being automated. Specifically, traditional, pre-AI types of automation, which primarily affected manual tasks (e.g., “performing detailed manual tasks” or “performing physical labor”) automate tasks for which we estimate a low skill dispersion. The advent of AI technologies, meanwhile, has brought about prospects of automation for those tasks for which skills are much more dispersed. We will come back to this discussion in Section 4, when we talk about the implications of various types of automation shocks.

[31] We take this to be a rather large value. This estimate arises because the data features substantial period-to-period wage variation even for individuals who stay in their occupation. However, this is of no consequence for our quantitative exercises below, since the realization of $\varepsilon_{i,t}$ does not have an impact on occupational choices or the relative occupation-specific wage prospects of workers.

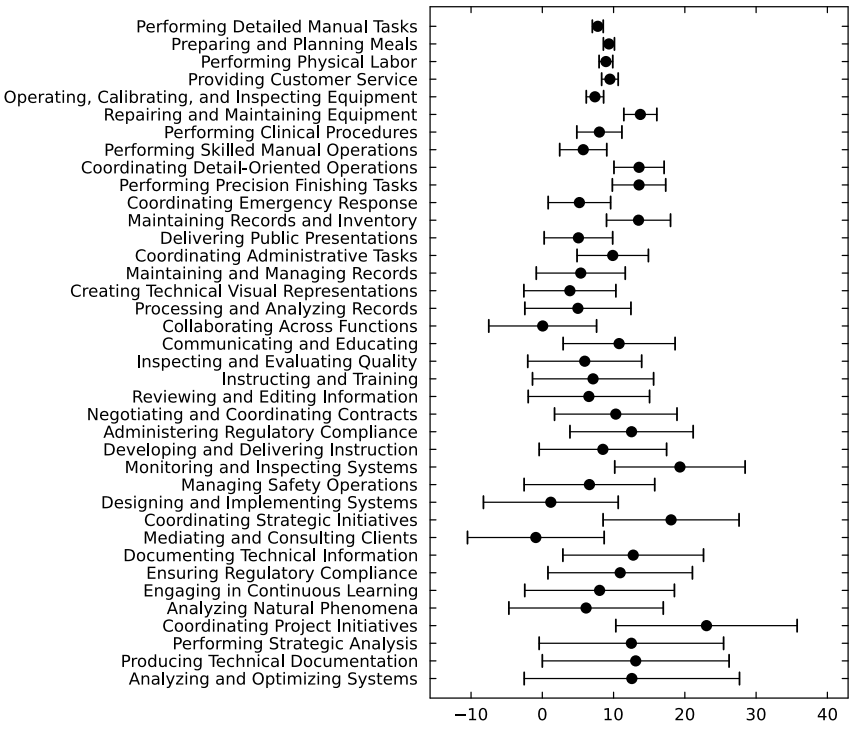


Figure 5: Estimated mean skills and dispersion

Notes: Dots show point estimates for mean log skills $\bar{s}_\tau$ by category τ. Plotted intervals cover one standard deviation of the corresponding marginal distribution in each direction. Tasks are ordered by increasing skill dispersion.

Lastly, we estimate the full correlation matrix C_s of all pairwise skills. This matrix is displayed in Appendix C.2.1.[32] In Appendix C.2.3, we also show that we obtain very similar parameter estimates when we use the NLSY97 in place of the NLSY79.

3.4.2 Model properties

We now describe the properties of the estimated model, starting with properties of the wage distribution.[33] In the data, the standard deviation of log wages is 0.60. Decomposing the total variance of log wages into within-occupation and between-occupation components, the within-term accounts for 71.5% of variation and the between term for 28.5%. The model does a solid job in matching these moments: The overall standard deviation is 0.70, with the within and between terms accounting for 18.6% and 81.4% of the total variance, respectively.

[32] It is difficult to directly validate the plausibility of any given correlational estimate. However, the Monte Carlo exercises conducted in Section 3.3 indicate that we are able to obtain a robust estimate of the correlation structure for any pair of tasks.

[33] All results are based on a simulation of 50,000 workers.

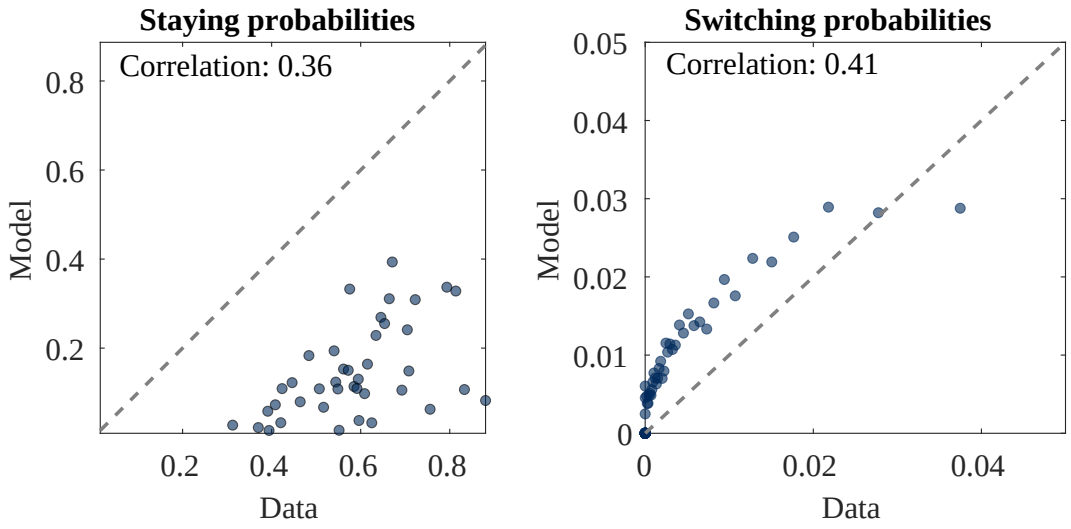


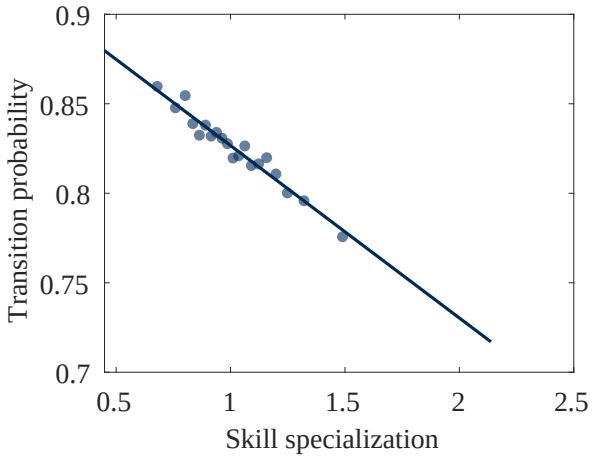
Figure 6: Occupational transition matrix: model vs. data

Notes. This figure compares the model-generated entries of the transition matrix to those derived from the NLSY. The left panel is a binned scatter plot of diagonal entries, the right panel a binned scatter of off-diagonal entries.

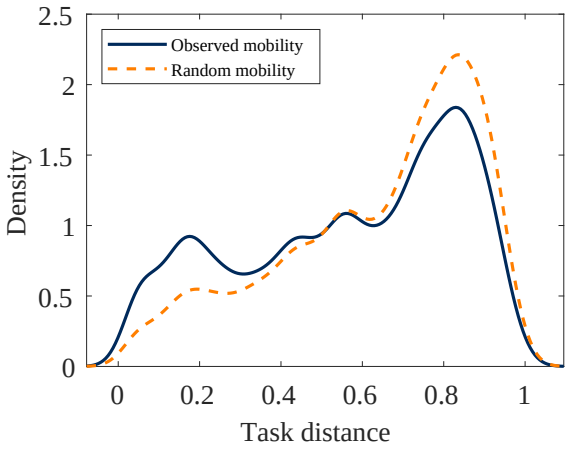
Turning to occupational choice, workers sort into occupations on the basis of task-level comparative advantage: In the model, workers sort into jobs that emphasize tasks where they possess relative skill advantages. For example, workers choosing to become Top Executives exhibit notably higher specialization in "'Coordinating Strategic Initiatives" – a task with a disproportionately high weight in this occupation – compared to the unconditional distribution of workers. This reflects the type of sorting captured in equation (12). Figure C.2 in the appendix provides an illustration.

What are the patterns of occupational switches in the estimated model, and how do they compare to the data? Figure 6 compares the model-implied occupational transition matrix entries to those observed in the NLSY. The model generates positive correlations with data in terms of both staying probabilities (diagonal elements, 0.36) and switching probabilities (off-diagonal elements, 0.41). While the model captures the directional pattern of occupational transitions, it notably under-predicts persistence in the same occupation: The average annual staying probability is 18%, well below the 63% measured in the NLSY. This gap likely reflects the absence of switching frictions or costs in our current specification, where transitions are driven purely by relative skill advantages and preference shocks.

An important piece of validation for our task-based model are its distinctive predictions for the effect of skill specialization on the *frequency* and the *direction* of occupational switches, which we show are at least qualitatively consistent with the data. First, the empirical evidence suggests that skill specialization tends to generate persistence in occupational choice (Kambourov and Manovskii, 2008; Geel *et al.*, 2011). Consistent with this idea, our theory implies that individuals



(a) More specialized workers move less



(b) Observed moves reflect task similarity

Figure 7: Skill specialization shapes the frequency and direction of moves

Notes. The left panel is a binscatter of individual-level observations, relating the normalized frequency of occupation switches to the coefficient of variation of skills. The right panel plots the observed density of distances conditional on switching occupation (solid line) and that under a random-mobility benchmark (dashed line).

with more specialized skills tend to move less. Our measure of individual skill specialization is the within-worker coefficient of variation of skills. Figure 7a shows that greater specialization is associated with a decreasing probability of switching occupation in any given period.

Next, consider the direction of moves conditional on occupational switching. Gathmann and Schönberg (2010) show, using German worker panel level data, that workers are more likely to move to occupations with similar task requirements. Whereas this relationship cannot be rationalized by models featuring skills that are either one-dimensional or fully occupation-specific, it is successfully reproduced by our model. To illustrate this property, we replicate the analysis conducted by Gathmann and Schönberg (2010) using data generated from our model. (Appendix Figure B.6 provides the empirical counterpart based on the NLSY.) Concretely, we compare the realized distribution of between-occupation distances traveled in the space of tasks to that we would observe if mobility was random, i.e., governed solely by the relative size of the destination occupation. The distance between any two occupations o and o' is one minus the angular separation of the task-weight vectors. Figure 7b demonstrates that under the observed distribution of distances more density is concentrated at shorter distances than under the random-mobility benchmark.

In summary, the estimated model, despite its simplicity, does a solid job at matching the wage distribution and captures important features of the relationship between task-specific skills and job mobility.

4 The Labor Market Effects of LLMs

We are now in a position to use the estimated model to construct projections of the wage effects of automation due to large language models (LLMs). Section 4.1 explains how we construct the automation shock. Section 4.2 evaluates occupation-level effects and decomposes them following the method derived in Section 2.3.3. Section 4.3 quantifies the consequences of automation at the individual level. Section 4.4 offers a brief discussion of the appropriate interpretation of our results.

4.1 Construction of automation shocks

This section explains how we leverage existing task exposure measures to identify which model tasks are affected by LLMs and details the specific automation scenario we analyze in the model.

Mapping to task exposure measures. Quantifying automation-induced earnings effects for technologies currently being rolled out, or to be adopted in the future, requires knowledge of which tasks are being, or will be, automated. This represents a non-trivial challenge. Unlike backward-looking studies, we cannot rely on labor share changes to measure automation in industries or occupations. Moreover, even if such changes could be constructed, they would not reveal which specific tasks within industries or occupations face automation. As highlighted in Section 2.3.3, this granular knowledge is essential for capturing job transformation effects.

Our methodology overcomes this challenge by providing a direct mapping between our model tasks—constructed from detailed O*NET tasks—and empirical measures of technology-specific, task-level automation exposure from a burgeoning literature using data sources ranging from patent data (Webb, 2019) and capability-specific AI benchmarks (Felten *et al.*, 2021) to expert and machine judgment (Eloundou *et al.*, 2023). Our model can link directly to any exposure measure at the detailed O*NET task level.

Motivated by the rapid diffusion of large language models (LLMs) with increasingly advanced capabilities (Bick *et al.*, 2024), we focus on this technology in our headline analysis. To identify which tasks are most likely automated through LLMs, we draw on Eloundou *et al.* (2023), who quantify O*NET task exposure to LLM automation using human labeling and GPT-4 classifications. We aggregate their scores to our task cluster level by averaging.[34] Figure 8 shows the resulting exposure scores for our task clusters, ordered by descending exposure. The most exposed categories are "Processing and Analyzing Records" and "Maintaining and Managing

[34] We focus on the automation rubric specifically, rather than the general exposure rubric. Eloundou *et al.* (2023) quantify exposure using five categorical bins. To convert these to numerical scores between 0 and 1, we use the authors' mapping: T0 (no automation exposure) receives a score of 0, T1 receives 0.25, T2 receives 0.5, T3 receives 0.75, and T4 (full automation exposure) receives 1.

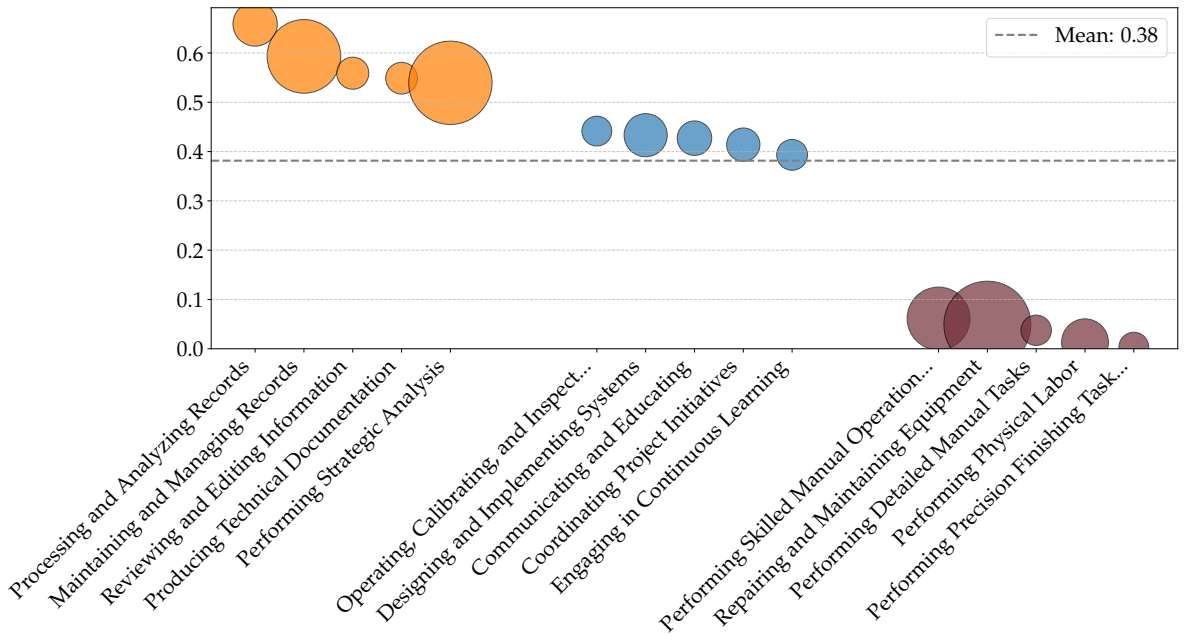


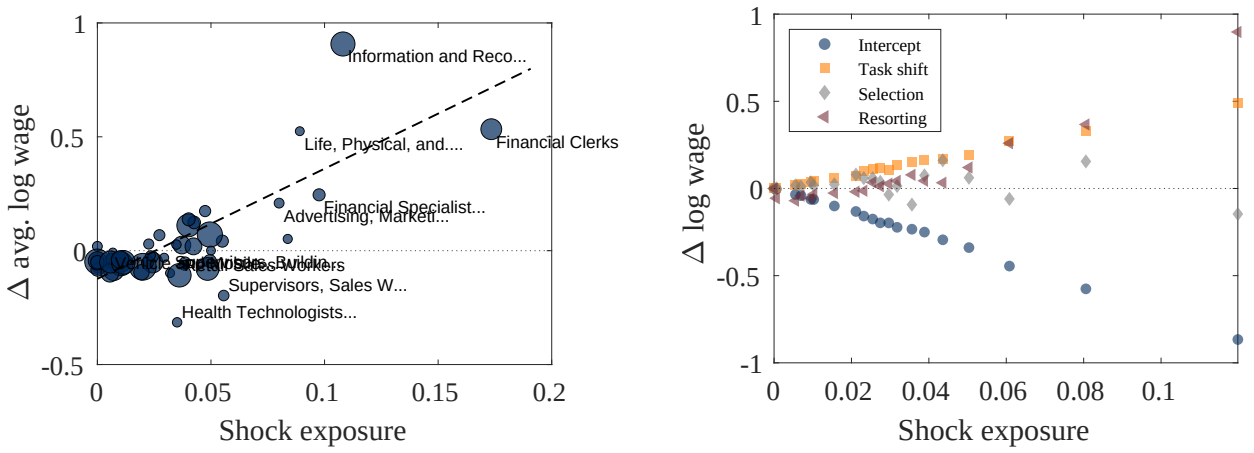
Figure 8: Eloundou *et al.* (2023) exposure scores aggregated to task clusters

Notes. This chart shows the exposure of the task clusters to LLM automation based on the aggregated exposure scores provided by Eloundou *et al.* (2023). The size of each bubble indicates the number of detailed O*NET tasks contained in each cluster. For example, the two most exposed task clusters are "Processing and Analyzing Records" (a cluster comprising 170 detailed tasks with average exposure of 0.66) and "Maintaining and Managing Records" (470 tasks with average exposure score 0.59).

Records," followed by "Reviewing and Editing Information" and "Producing Technical Documentation." The first two clusters include detailed, occupation-specific tasks such as "Compute payment schedules" and "Prepare reports showing places of departure and destination, passenger ticket numbers, [...]" or "Maintain and update human resources documents, such as organizational charts, [...]" and "Organize archival records and develop classification systems to facilitate access to archival materials," respectively.

To contextualize the model-implied labor market consequences of LLMs, we compare them to those from automation by industrial robots. Our approach's flexibility allows us to measure robot automation exposure for our task clusters by linking them to Webb (2019) exposure measures. Webb (2019) construct task automation exposure measures using the overlap between job task descriptions and patent text.[35] This methodology identifies "Performing Detailed Manual Tasks" as the most robot-exposed task cluster—comprising detailed occupation-specific tasks such as "Lubricate moving parts" and "Remove excess materials or impurities from objects, using air hoses or grinding machines"—followed by "Performing Physical Labor," which includes tasks like "Hammer out bulges or bends in metal workpieces" and "Dump refuse or recyclable materials at disposal sites." Appendix B.5 provides more details.

[35] Many thanks to Mike Webb for sharing the exposure scores at the task level.



(a) Effects **(b)** Decomposition

Figure 9: LLM automation – occupation-level effects

Notes. Left panel: Each dot is an occupation. The vertical axis measures the average occupational wage change before versus after the shock. The horizontal axis measures shock exposure $A_{o,\tau}$. Dot sizes correspond to pre-shock employment shares. The dotted line is the line of best fit. Right panel: Bin scatter of occupational effect size according to equation (11) by shock exposure $A_{o,\tau}$. Note that the horizontal axis range differs across the two panels due to binning in the right panel.

Automation scenario. In the following sections, we use the model to evaluate the earnings effects when the task most exposed to LLMs, i.e., "Processing and Analyzing Records" is fully automated. Appendix C.2.2 shows that results are similar when considering the closely related task "Maintaining and Managing Records" – sometimes we refer to them jointly as information-processing tasks. Throughout, we assume that machine productivity in the automated task, z_{τ^*}, is at its automation threshold, defined in Section 2.3.1. This means that z_{τ^*} is *just* high enough to make the automation of task $\tau^\star$ optimal, so the scenario we are considering can be interpreted as a lower bound on average productivity and wage effects.

4.2 Occupation-level effects & their limits

We begin by studying what happens to occupation-level wages when "Processing and Analyzing Records" is fully automated. Panel (a) of Figure 9 plots occupational wage changes from automating this task against occupational exposure, measured by $A_{o,\tau}$. The figure shows that occupational wage effects tend to be positive. More exposed occupations experience larger wage gains on average. Occupations with the highest loadings on the automated task, such as Financial Clerks and Information and Record Clerks, experience the largest increases in average occupational wages.

The decomposition in equation 11 provides a natural framework to explain these effects. Panel (b) of Figure 9 presents the estimated decomposition terms as a bin scatter, plotting the

average magnitude of each effect against occupational exposure. The figure shows that these positive effects are not driven by traditional displacement and productivity effects captured by $\Delta\mu_o$; indeed, these terms tend to be more negative for more exposed occupations. This negative effect, however, is offset by three job transformation effects: task shifting, selection, and re-sorting.

First, automating "Processing and Analyzing Records" leads to *task upgrading*, indicated by positive task shift effects. Automation frees up production time and allows workers to spend more time on tasks where the average worker is more productive. This increases productivity and average occupational wages. The effect is stronger for more exposed occupations, where more time is freed up by automation.

Second, selection effects have ambiguous effects on average wages as exposure grows. This shows that selection plays a subtle role: For "Processing and Analyzing Records," workers select into exposed occupations based on the automated skill but also on skills other than the automated one. When automation occurs, selection effects are thus ambiguous, hurting some workers while benefiting others. Selection effects benefit workers who initially select occupations based on non-automated skills but harm workers who select mainly on automated skills.[36]

Third, re-sorting effects drive the average wage gains for the most affected occupations. This reflects that the shock generates significant labor market "turbulence." The composition of workers in the most exposed occupations changes substantially as automation hits: existing workers leave and new workers enter as the type of work changes.

The magnitude of re-sorting among the most exposed occupations suggests that occupational averages may lack economic meaning and provide poor guidance for understanding *individual-level* effects. We therefore ask under what conditions re-sorting effects are large. To understand this, we compare all 38 possible full automation shocks by the slope they induce in their respective versions of Figure 9. For each shock, we examine how much more exposed occupations experience larger re-sorting effects. Figure 10 plots these measures against the dispersion of the associated skill. Orange highlights the two tasks most exposed to AI in our aggregation of Eloundou *et al.* (2023). Gray highlights the tasks most exposed to robots in our aggregation of Webb (2019). Figure 10 reveals a clear pattern: shocks affecting highly dispersed skills generate larger re-sorting effects; shocks affecting less dispersed skills produce small re-sorting effects. Equation (11) explains why: shocks to highly dispersed skills feature larger re-sorting effects when negatively correlated with other high-dispersion skills strongly utilized by affected occupations. This holds for our estimated parameters.[37]

[36] Appendix C.2 shows this result is somewhat unique to this task: other automation types can generate much more pronounced negative wage effects for most incumbents. Across shocks, selection effects rarely generate wage gains in affected occupations.

[37] Appendix C.2 discusses how automation shock properties also affect selection and task-shifting effects.

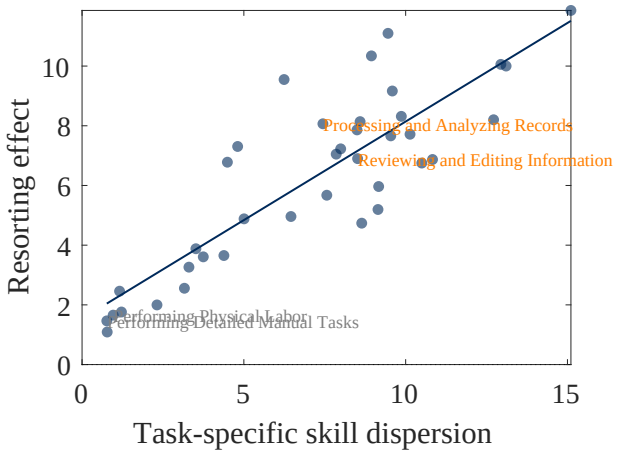


Figure 10: Resorting effect – comparison across tasks

Notes. Each dot corresponds to a task. The horizontal axis displays the estimated between-worker dispersion of the skill at task τ: $S_{s,\tau}$. The vertical axis displays the coefficient of a regression that regresses the magnitude of the re-sorting effect from full automation for a given occupation on its exposure $A_{o,\tau}$, weighting coeffcients by pre-shock occupational employment.

Importantly, our estimation suggests AI-exposed tasks ("Processing and Analyzing Records," "Reviewing and Editing Information") tend to be associated with larger skills dispersion than is the case for tasks automated by industrial robots ("Performing Detailed Manual Tasks," "Performing Physical Labor"). Occupation-level averages therefore provide worse guidance for worker-level outcomes under AI-driven automation than under historical robot automation.

In sum, AI-driven automation leads to positive average wage effects for the most exposed occupations, driven by task upgrading and a changing pool of workers. The next section turns to an analysis of individual-level effects. While these effects do not admit elegant closed-form characterization, many insights from equation (11) carry over.

4.3 Individual-level effects

Incumbent workers in exposed occupations experience, on average, wage losses, contrasting sharply with positive occupation-level effects. To show this, we classify individuals by their origin occupation before the shock and track their earnings over time. Figure 11 plots average wage effects for all incumbent workers by their origin occupation against occupational exposure. The figure reveals that whereas occupation-level wages rise with exposure (Figure 9), incumbent wages fall with exposure.

This contrast arises because of shifts in occupational worker composition. To see this, observe

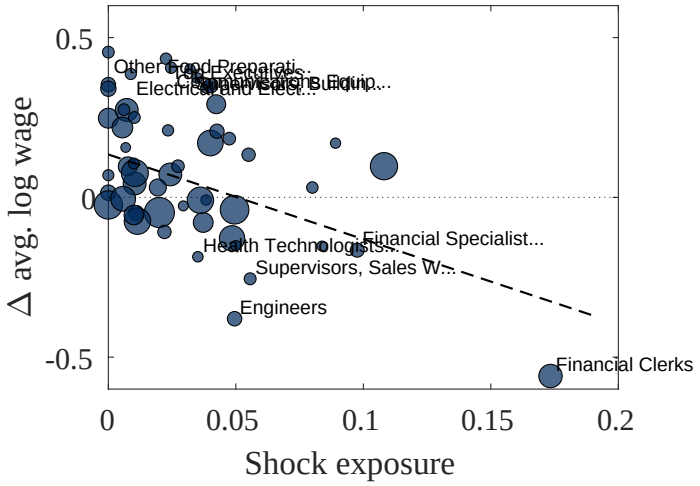


Figure 11: LLM automation – individual-level wage effects for incumbents

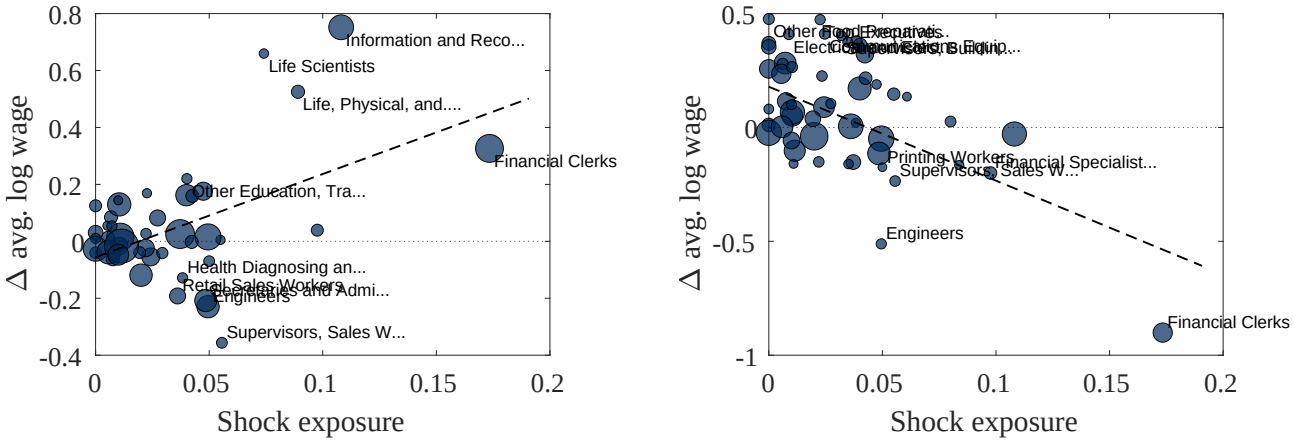
Notes. Each dot is an occupation. The vertical axis measures the wage change for initial incumbents of each occupation before versus after the shock. The horizontal axis measures shock exposure $A_{o,\tau}$. Dot sizes correspond to pre-shock employment shares. The dotted line is the line of best fit.

that incumbent wage effects can be decomposed as follows:

$$\begin{aligned}&\sum_{o''\in O} \lambda'_{o''}(s_{|o})\mathbb{E}[w'_{o''}|\hat{o}=o] - \mathbb{E}[w_o|\hat{o}=o] \\ &= \underbrace{\mathbb{E}[w'_o|\hat{o}=o] - \mathbb{E}[w_o|\hat{o}=o]}_{\Delta w_o \text{ of incumbents}} + \underbrace{\sum_{o''\in O} \lambda'_{o''}(s_{|o})\mathbb{E}[w'_{o''} - w'_o|\hat{o}=o]}_{\text{Reallocation of incumbents}}. \end{aligned} \tag{14}$$

where $\lambda_{o''}(s_{|o})$ denotes the post-automation employment share of occupation o'' for workers initially in o. The first term can be decomposed according to equation (11) into productivity and displacement, task-shift, and selection effects. Since equation (14) describes incumbent wage effects, the "re-sorting" term from equation (11) becomes a term describing incumbent reallocation across occupations. The contrast between occupational averages and the experience of individual incumbents thus arises from two sources: first, wage changes at the occupational level are partly driven by non-incumbents; second, incumbents themselves may move to other jobs with wages that differ from those in their origin occupation. Figure 11 indicates that incumbents in exposed occupations do worse than their occupational average suggests. Moreover, their reallocation to new occupations does not successfully insure them against wage losses.

The importance of reallocation motivates us to distinguish between those incumbents who decide to stay in their job and those who move into other occupations — a distinction we show is associated with large differences in outcomes. Figure 12 splits the population of incumbents into these two groups. The figure reveals substantial heterogeneity within occupational incumbents:



(a) Incumbent stayers **(b)** Incumbent leavers

Figure 12: LLM automation – stayers versus leavers

Notes. Each dot is an occupation. The vertical axis of panel (a) measures the wage change for initial incumbents of each occupation before versus after the shock, including only those who stay in their occupation after the shock. The vertical axis panel (b) measures the wage change for initial incumbents of each occupation before versus after the shock, including only those who leave their occupation after the shock. The horizontal axis of both panels measures shock exposure $A_{o,\tau}$. Dot sizes correspond to pre-shock employment shares. The dashed line is the line of best fit.

stayers do much better than leavers. Figure 13 explains why, highlighting divergent patterns of skill specialization. The figure plots relative specialization, defined as $s_{o,\tau^\star} - \frac{1}{n_{\text{skill}}} \sum_{\tau \in \mathcal{T}_l} s_{o,\tau}$, for incumbents in low- and high-exposure occupations and, within the latter group, incumbent stayers and incumbent leavers. This contrast shows that selection plays a major role in generating differences between stayers and leavers: leavers are highly specialized in automated tasks, whereas the degree of specialization for stayers is similar to that of the general population.

The degree of dispersion in wage effects among incumbents workers can be traced back to the properties of the skill distribution. Focusing on origin-occupation wages allows the following characterization:

$$\text{Var}(\Delta w_{i,o} | \hat{o} = o) = (A'_o - A_o)\Sigma_{s|o}(A'_o - A_o)^\intercal$$

$$= (A'_o - A_o)\left(\Sigma_s^{-1} + \nu^{-2} \underbrace{\left(\sum_{o'} h_{o'}(\bar{s}_{|o}) A_{o'}^\intercal A_{o'} - \left(\sum_{o'} h_{o'}(\bar{s}_{|o}) A_{o'}^\intercal \right) \left(\sum_{o'} h_{o'}(\bar{s}_{|o}) A_{o'}^\intercal \right)^\intercal \right)}_{\text{Task intensity dispersion across occupations}} \right)^{-1} (A'_o - A_o)^\intercal.$$

This shows that the automation of tasks associated with greater skill dispersion tends to lead to more dispersion among incumbents, especially when occupations do not differ much in their

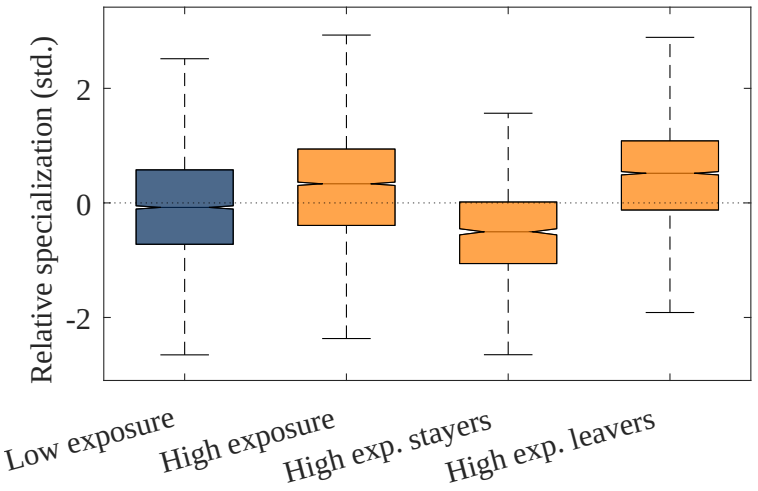


Figure 13: LLM automation – specialization in automated task by worker groups

Notes. Each box plot shows the distribution of $s_{i,\tau^\star} - \frac{1}{n_{\text{skill}}} \sum_{\tau \in \mathcal{T}_l} s_{i,\tau}$ where $\tau^\star$ corresponds to the "Processing and Analyzing Records" task. "Low exposure" ("high exposure") refers to those workers who are in the bottom (top) 10% by exposure, as measured by $A_{o,\tau^\star}$ of their occupation.

exposure and, thus, initial sorting is less pronounced.

To appreciate who the winners and losers are following LLM-driven automation, consider a few examples. Workers whose skills are concentrated in information-processing tasks lose, exiting incumbent jobs that no longer reward their comparative advantage. In contrast, many other workers, especially in highly exposed office and administrative roles, stay and gain as their work content shifts. This includes, for instance, workers who stay in occupations such as "Financial Clerks," which place large weights on customer-facing and coordination tasks. Similarly, lawyers who remain in their occupation and experience wage gains are typically those whose skills concentrate in communication ("Communicating and Educating") and negotiation ("Negotiating and Coordinating Contracts") rather than document processing and analysis.

These model-based findings are consistent with heterogeneous adjustment patterns documented in the empirical literature. Thus, Dauth *et al.* (2021) document, using German micro data, that workers who stayed at their firms after the advent of industrial robots tended to experience wage gains and shifted their work toward more productive tasks, whereas those who left tended to lose. Similarly, Gathmann *et al.* (2024) provide suggestive evidence that automation of information-processing tasks through AI benefits a subset of workers who are able to reallocate toward other tasks in which they are proficient. Additionally, Huckfeldt (2022) shows that, in general, the earnings cost of job loss is almost entirely concentrated among workers who subsequently switch occupation.

Incumbent stayers who benefit from task upgrading are not the only winners. As the large re-sorting effect suggests, many workers realize wage gains because automating a task removes

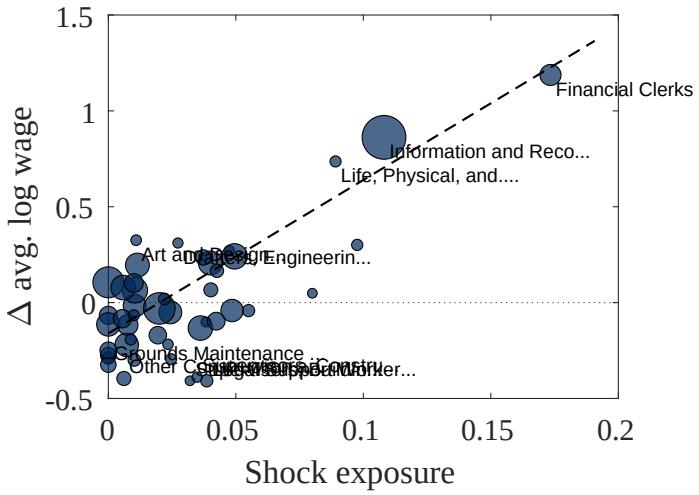


Figure 14: LLM automation – in-switcher wage effects

Notes. Each dot is an occupation. For each occupation, the vertical axis measures the wage change of workers who are *not* incumbents of that occupation but select that occupation after the shock. The horizontal axis measures shock exposure $A_{o,\tau}$. Dot sizes correspond to pre-shock employment shares. The dotted line is the line of best fit.

entry barriers into occupations they would otherwise qualify for. Figure 14 illustrates this point by comparing wage effects of workers who switch into each occupation against occupational exposure. In-switchers into the most exposed occupations benefit the most. This reflects that these occupations undergo substantial changes in task content and thus attract workers whose skills fit better with the new profile of the transformed occupation.[38]

Together, these findings paint a picture of diverging fortunes among individuals in exposed occupations. On the one hand, automation creates losers: Those particularly skilled in the automated task who selected into the most exposed occupations precisely because of this specialization. These workers tend to leave after the shock but struggle to find alternative employment at or above their previous wage level. On the other hand, automation also creates winners: workers who selected exposed occupations for reasons other than their skill at the automated task. These workers benefit from task upgrading, typically stay after automation, and realize wage gains. A final group of winners are those who initially select less exposed occupations but who benefit from the shift in task content and after the shock are able to gainfully move into the most affected occupations and utilize their comparative advantages.

[38] Going beyond the specific automation scenario we consider here, an evocative but speculative example of this effect may be the rise of "vibe coding," which may allow individuals who previously would never have worked as software engineers due to a lack of coding skills to enter the profession, as LLM code generation shifts the job toward higher-level tasks like creativity as well as project planning and quality control."

4.4 Discussion

We close by discussing the appropriate interpretation of our quantitative results and potential extensions. As job transformation effects are typically abstracted from in quantitative automation analyses, we made three assumptions to transparently identify these effects. First, we deliberately formulated our model in partial equilibrium. This approach demonstrates that a distinct set of distributional effects arises even when occupational prices are fixed. Introducing general equilibrium effects could dampen or amplify their magnitude.

Second, the model features no exogenous switching costs or frictions. This assumption traces any non-random job mobility to skill heterogeneity, but results in limited occupational persistence which may translate into overstated wage gains accruing to in-switchers. In addition, to facilitate estimation and analytical characterization, we assumed skills to be time-invariant. As mentioned previously, this means our results are best interpreted as applying to a horizon of three to five years following the shock. Appendix C.2.5 sketches an extension of the model allowing for a simple form of learning in the form of returns to occupational experience. We show this extension allows the model to generate greater occupational persistence.

Third, we assume no heterogeneity along observable dimensions such as age or gender to isolate the effects of skill heterogeneity. It would be valuable for future work to explore how job transformation effects vary with demographic characteristics.

5 Conclusion

If the historical record – and anecdotal evidence about work reorganization at firms adopting gen-AI at large scale – offers any guidance, job task transformation will play a first-order role in shaping the labor market effects of AI. We develop a framework to analyze these effects, offering both conceptual insights and methodological tools to quantify them.

Our findings underscore that automation can be disruptive for workers even it does not *eliminate* jobs. In our model, AI automation generates large and heterogeneous wage effects by *transforming* the task content of jobs. Furthermore, our analysis challenges the common practice of equating task- or occupation-level automation exposure with negative wage effects. Conceptually, job transformation effects create heterogeneous effects at the individual level – workers exposed to the same shock in the same occupation may fare very differently depending on their relative specialization.[39] Quantitatively, our findings underscore that occupation-level averages are especially misleading for AI-induced task automation, because AI automates tasks

[39] In addition, greater exposure can be positive or negative depending on the relative magnitude of displacement and productivity effects. In our model, positive productivity effects also accrue only to exposed occupations.

with larger associated skill dispersion than in past episodes, prompting larger shifts in occupational skill composition. Incumbents in the same occupation include winners, who stay and move to higher-productivity tasks, and losers, who are pushed out into lower-paying jobs. In brief, exposure measures provide valuable insights into which tasks specific technologies affect, but deriving implications for earnings requires a structural model that carefully maps exposure into wages.

Our framework offers several attractive properties that, we hope, will be of value for other researchers. First, because it links directly to task exposure measures, including forward-looking ones, it can analyze labor market consequences without waiting for backward-looking exposure data, which is especially useful for policy. Second, as the model tracks individuals' earnings, not only occupational outcomes, the framework can be used to evaluate whether AI will exacerbate wage inequality, as has been argued for past automation shocks (Acemoglu and Restrepo, 2022), or, as has been conjectured, might in fact dampen it.[40] Third, our counterfactual analyses allow for arbitrarily large shocks with potentially non-linear effects. Fourth, data requirements are limited: worker-level panel data are widely available, and we offer a flexible methodology to measure occupational task weights. Thus, the framework can readily be used to extend the quantitative analysis of job transformation effects to more countries.

[40] See, for instance, Webb (2019) or David Autor in https://www.noemamag.com/how-ai-could-help-rebuild-the-middle-class.

References

Acemoglu, D. and Autor, D. (2011). Skills, Tasks and Technologies: Implications for Employment and Earnings. In *Handbook of Labor Economics*, volume 4, pages 1043–1171. Elsevier.

Acemoglu, D. and Restrepo, P. (2018a). Artificial Intelligence, Automation and Work.

Acemoglu, D. and Restrepo, P. (2018b). The Race between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment. *American Economic Review*, **108**(6), 1488–1542.

Acemoglu, D. and Restrepo, P. (2022). Tasks, Automation, and the Rise in U.S. Wage Inequality. *Econometrica*, **90**(5), 1973–2016.

Acemoglu, D., Kong, F., and Restrepo, P. (2025). Tasks At Work: Comparative Advantage, Technology and Labor Demand. *mimeo*.

Aghion, P., Jones, B. F., and Jones, C. I. (2017). Artificial Intelligence and Economic Growth.

Althoff, L. and Reichardt, H. (2025). AI and Comparative Advantage. *mimeo*.

Atalay, E., Phongthiengtham, P., Sotelo, S., and Tannenbaum, D. (2020). The Evolution of Work in the United States. *American Economic Journal: Applied Economics*, **12**(2), 1–34.

Athey, S., Brunborg, H., Du, T., Kanodia, A., and Vafa, K. (2024). LABOR-LLM: Language-Based Occupational Representations with Large Language Models.

Autor, D. and Thompson, N. (2025). Expertise.

Autor, D., Chin, C., Salomons, A., and Seegmiller, B. (2024). New Frontiers: The Origins and Content of New Work, 1940–2018. *The Quarterly Journal of Economics*, **139**(3), 1399–1465.

Autor, D. H. (2015). Why Are There Still So Many Jobs? The History and Future of Workplace Automation. *Journal of Economic Perspectives*, **29**(3), 3–30.

Autor, D. H. and Dorn, D. (2013). The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market. *American Economic Review*, **103**(5), 1553–1597.

Autor, D. H. and Handel, M. J. (2013). Putting Tasks to the Test: Human Capital, Job Tasks, and Wages. *Journal of Labor Economics*, **31**(2), S59–S96.

Autor, D. H., Levy, F., and Murnane, R. J. (2003). The Skill Content of Recent Technological Change: An Empirical Exploration. *The Quarterly Journal of Economics*, **118**(4), 1279–1333.

Baley, I., Figueiredo, A., and Ulbricht, R. (2022). Mismatch Cycles. *Journal of Political Economy*, **130**(11), 2943–2984.

Bartel, A., Ichniowski, C., and Shaw, K. (2007). How Does Information Technology Affect Productivity? Plant-Level Comparisons of Product Innovation, Process Improvement, and Worker Skills. *The Quarterly Journal of Economics*, **122**(4), 1721–1758.

Bessen, J. (2012). More Machines, Better Machines... or Better Workers? *The Journal of Economic History*, **72**(1), 44–74.

Bessen, J. E. (2011). Was Mechanization De-Skilling? The Origins of Task-Biased Technical Change. *SSRN Electronic Journal*.

Bick, A., Blandin, A., and Deming, D. J. (2024). The Rapid Adoption of Generative AI.

Bocquet, L. (2022). The Network Origin of Slow Labor Reallocation. *Working Papers*, (halshs-03703862).

Böhm, M., Etheridge, B., and Irastorza-Fadrique, A. (2025). The Impact of Labour Demand Shocks When Occupational Labour Supplies are Heterogeneous.

Bonney, K., Breaux, C., Buffington, C., Dinlersoz, E., Foster, L., Goldschlag, N., Haltiwanger, J., Kroff, Z., and Savage, K. (2024). The impact of AI on the workforce: Tasks versus jobs? *Economics Letters*, **244**, 111971.

Boustan, L. P., Choi, J., and Clingingsmith, D. (2022). Computerized Machine Tools and the Transformation of US Manufacturing.

Brynjolfsson, E., Mitchell, T., and Rock, D. (2018). What Can Machines Learn, and What Does It Mean for Occupations and the Economy? *AEA Papers and Proceedings*, **108**, 43–47.

Brynjolfsson, E., Li, D., and Raymond, L. (2025). Generative AI at Work*. *The Quarterly Journal of Economics*, **140**(2), 889–942.

Caselli, F. and Manning, A. (2019). Robot Arithmetic: New Technology and Wages. *American Economic Review: Insights*, **1**(1), 1–12.

Cazzaniga, M. (2024). Gen-AI. *IMF Staff Discussion Notes*, **2024**(001), 1.

Chequer, M., Herkenhoff, K., Papanikolaou, D., Schmidt, L., and Seegmiller, B. (2025). An Eddie Lazear Model of AI and Labor Markets: Theory Meets Resume Data. *mimeo*.

Dauth, W., Findeisen, S., Suedekum, J., and Woessner, N. (2021). The Adjustment of Labor Markets to Robots. *Journal of the European Economic Association*, **19**(6), 3104–3153.

del Rio-Chanona, R. M., Mealy, P., Beguerisse-Díaz, M., Lafond, F., and Farmer, J. D. (2021). Occupational mobility and automation: A data-driven network model. *Journal of The Royal Society Interface*, **18**(174), 20200898.

Dell, M. (2024). Deep Learning for Economists.

Deming, D. J. (2023). Multidimensional Human Capital and the Wage Structure.

Dix-Carneiro, R. (2014). Trade Liberalization and Labor Market Dynamics. *Econometrica*, **82**(3), 825–885.

Eloundou, T., Manning, S., Mishkin, P., and Rock, D. (2023). GPTs are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models.

Fan, T. (2025). The Labor Market Incidence of New Technologies.

Fan, T. and Restrepo, P. (2025). Partial Automation. *mimeo*.

Felten, E., Raj, M., and Seamans, R. (2021). Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses. *Strategic Management Journal*, **42**(12), 2195–2217.

Felten, E. W., Raj, M., and Seamans, R. (2018). A Method to Link Advances in Artificial Intelligence to Occupational Abilities. *AEA Papers and Proceedings*, **108**, 54–57.

Freund, L. B. (2023). Superstar Teams. *JIWP Number: 2235*.

Frey, C. B. and Osborne, M. A. (2017). The future of employment: How susceptible are jobs to computerisation? *Technological Forecasting and Social Change*, **114**, 254–280.

Garicano, L. (2000). Hierarchies and the Organization of Knowledge in Production. *Journal of Political Economy*, **108**(5), 874–904.

Gathmann, C. and Schönberg, U. (2010). How General Is Human Capital? A Task-Based Approach. *Journal of Labor Economics*, **28**(1), 1–49.

Gathmann, C., Grimm, F., and Winkler, E. (2024). Ai, Task Changes in Jobs, and Worker Reallocation.

Geel, R., Mure, J., and Backes-Gellner, U. (2011). Specificity of occupational training and occupational mobility: An empirical study based on Lazear's skill-weights approach. *Education Economics*, **19**(5), 519–535.

Grigsby, J. (2023). Skill Heterogeneity and Aggregate Labor Market Dynamics. *mimeo*.

Grigsby, J. and Zorzi, N. (2025). The Labor Market Consequences of a Rapid Climate Transition. *mimeo*.

Guvenen, F., Kuruscu, B., Tanaka, S., and Wiczer, D. (2020). Multidimensional Skill Mismatch. *American Economic Journal: Macroeconomics*, **12**(1), 210–244.

Hampole, M., Papanikolaou, D., Schmidt, L. D., and Seegmiller, B. (2025). Artificial Intelligence and the Labor Market.

Hsieh, C.-T., Hurst, E., Jones, C. I., and Klenow, P. J. (2019). The Allocation of Talent and U.S. Economic Growth. *Econometrica*, **87**(5), 1439–1474.

Huckfeldt, C. (2022). Understanding the Scarring Effect of Recessions. *American Economic Review*, **112**(4), 1273–1310.

Humlum, A. (2019). Robot adoption and labor market dynamics. *mimeo*.

Humlum, A. and Vestergaard, E. (2025a). Large Language Models, Small Labor Market Effects.

Humlum, A. and Vestergaard, E. (2025b). The unequal adoption of ChatGPT exacerbates existing inequalities among workers. *Proceedings of the National Academy of Sciences*, **122**(1), e2414972121.

Ide, E. and Talamas, E. (2025). Artificial Intelligence in the Knowledge Economy. *Journal of Political Economy*.

Jones, C. I. (2022). The Past and Future of Economic Growth: A Semi-Endogenous Perspective. *Annual Review of Economics*, **14**(Volume 14, 2022), 125–152.

Kambourov, G. and Manovskii, I. (2008). Rising Occupational and Industry Mobility in the United States: 1968-97. *International Economic Review*, **49**(1), 41–79.

Kogan, L., Papanikolaou, D., Schmidt, L. D., and Seegmiller, B. (2023). Technology and Labor Displacement: Evidence from Linking Patents with Worker-Level Data.

Lashkari, D., Qiu, C., Li, W., and Thompson, N. (2025). AI, Scale, and Skills: A Quantitative Task-Based Theory of Automation. *mimeo*.

Lazear, E. P. (2009). Firm-Specific Human Capital: A Skill-Weights Approach. *Journal of Political Economy*, **117**(5), 914–940.

Lin, J. (2011). Technological Adaption, Cities, and New Work. *The Review of Economics and Statistics*, page 21.

Lindenlaub, I. (2017). Sorting Multidimensional Types: Theory and Application. *The Review of Economic Studies*, **84**(2), 718–789.

Lise, J. and Postel-Vinay, F. (2020). Multidimensional Skills, Sorting, and Human Capital Accumulation. *American Economic Review*, **110**(8), 2328–2376.

McInnes, L., Healy, J., and Astels, S. (2017). Hdbscan: Hierarchical density based clustering. *Journal of Open Source Software*, **2**(11), 205.

Moll, B., Rachel, L., and Restrepo, P. (2022). Uneven Growth: Automation's Impact on Income and Wealth Inequality. *Econometrica*, **90**(6), 2645–2683.

Noy, S. and Zhang, W. (2023). Experimental evidence on the productivity effects of generative artificial intelligence. *Science*, **381**(6654), 187–192.

Ocampo Díaz, S. (2022). A task-based theory of occupations with multidimensional heterogeneity. Technical report, mimeo.

Restrepo, P. (2024). Automation: Theory, Evidence, and Outlook. *Annual Review of Economics*, **16**(Volume 16, 2024), 1–25.

Roy, A. D. (1951). Some Thoughts on the Distribution of Earnings. *Oxford Economic Papers*, **3**(2), 135–146.

Sanders, C. (2012). Skill Uncertainty, Skill Accumulation, and Occupational Choice. *2012 Meeting Papers*, (633).

Spitz-Oener, A. (2006). Technical Change, Job Tasks, and Rising Educational Demands: Looking outside the Wage Structure. *Journal of Labor Economics*, **24**(2), 235–270.

Susskind, D. (2020). *A World Without Work: Technology, Automation, and How We Should Respond*. Metropolitan Books, New York, N.Y.

Traiberman, S. (2019). Occupations and Import Competition: Evidence from Denmark. *American Economic Review*, **109**(12), 4260–4301.

Webb, M. (2019). The impact of artificial intelligence on the labor market. *Available at SSRN 3482150*.

Woessmann, L. (2024). Skills and Earnings: A Multidimensional Perspective on Human Capital.

Online Appendix

This appendix contains supplemental material. Any references to sections, equations, figures, or tables that are not preceded by a capital letter refer to the main article.

A Theory appendix

A.1 Endogenizing $(\mathcal{T}_l, \mathcal{T}_m)$ & the automation threshold $\bar{z}_{\tau^\star}$

Section 2.1 treats the assignment of production tasks to labor and machines, $(\mathcal{T}_l, \mathcal{T}_m)$, as exogenous. We now discuss a set of additional assumptions that allow us to endogenize these sets as firms' choices. This allows us to determine, for any task $\tau^\star$ an "automation threshold" $\bar{z}_{\tau^\star}$ that triggers the optimal automation of this task. In what follows, we spell out these assumptions.

Entrepreneurs. There is a large mass of entrepreneurs. In every period, every worker randomly matches with $N \geq 2$ entrepreneurs. Before the occupation and skill are revealed to the entrepreneur, the entrepreneur makes an automation decision. That is, they decide the set of tasks that are produced with human labor, $\mathcal{T}_l$, and the set of tasks done by machines, $\mathcal{T}_m$. After automation decisions are taken, the occupation o and the worker's characteristics $(s_{i,\cdot}, \varepsilon_{i,t})$ are revealed. Wages are then set via Bertrand competition. Lastly, the winning entrepreneur forms a match with the worker and optimally allocates the worker's time to human tasks and machine capital to machine tasks.

Automation choice. Given some vector $\{z_\tau\}_{\tau\in\mathcal{T}}$, we define an optimal automation choice as task sets $(\mathcal{T}_l, \mathcal{T}_m)$ such that no entrepreneur finds it optimal to deviate from this task assignment. Note that the wage paid to a given worker is independent of the automation choice from the perspective of an individual firm considering a deviation. Thus, for any task τ, the condition that no firm finds it optimal to deviate from the assignment $(\mathcal{T}_l, \mathcal{T}_m)$ can be written as

$$\begin{aligned}&\int \left(\max_{m'} y'(m') - \exp(w(s,o,\varepsilon)) - rm'\right) dF(s|o)dG(\varepsilon)d\Lambda(o)\\ &\leq \int \left(\max_{m} y(m) - \exp(w(s,o,\varepsilon)) - rm\right) dF(s|o)dG(\varepsilon)d\Lambda(o) \qquad \text{(A.1)}\end{aligned}$$

where y' denotes the production function under a given alternative choice of task sets $(\mathcal{T}_l', \mathcal{T}_m')$,

and Λ, G, and F denote the distributions of occupational choices, idiosyncratic shocks ε (which are independent of occupational choices), and skills s conditional on occupational choices, respectively. The task assignment $(\mathcal{T}_l, \mathcal{T}_m)$ is thus optimal if and only if, for any alternative task assignments $(\mathcal{T}_l', \mathcal{T}_m')$, equation (A.1) is satisfied.

Using the expectations operator in place of integrals and substituting in optimality conditions, we can also write equation (A.1) as

$$\begin{aligned}
&\mathbb{E}_{(s|o),\varepsilon,o}\left[\exp\left(\mu_o' + \sum_{\tau\in\mathcal{T}_l'} \frac{\alpha_{o,\tau}}{\sum_{\tau\in\mathcal{T}_l'}\alpha_{o,\tau}} s_{i,\tau} + \varepsilon_{i,t}\right)\right] \\
&\leq \mathbb{E}_{(s|o),\varepsilon,o}\left[\exp\left(\mu_o + \sum_{\tau\in\mathcal{T}_l} \frac{\alpha_{o,\tau}}{\sum_{\tau\in\mathcal{T}_l}\alpha_{o,\tau}} s_{i,\tau} + \varepsilon_{i,t}\right)\right].
\end{aligned} \tag{A.2}$$

This establishes that the optimal automation threshold is the one that would leave average wages constant if occupational choices were held constant.

It can be shown that it is always possible to find values of $\{z_\tau\}_{\tau\in\mathcal{T}}$ that justify a given initial task assignment $(\mathcal{T}_l, \mathcal{T}_m)$ as optimal. For each task τ, we define the automation threshold $\bar{z}_\tau$ as the point at which Equation (A.2) holds with equality. That is, holding occupational choices constant, the average wage in the economy stays constant whether or not task τ is automated. It can be verified that, given an initially optimal assignment, there is such a threshold value $\bar{z}_\tau$ and it is finite for any task.

Equilibrium with endogenous automation. An equilibrium with endogenous automation is defined as a tuple of automation choices $(\mathcal{T}_l, \mathcal{T}_m)$ and a joint distribution Γ of occupation choices, log wages w, log skills s and idiosyncratic productivity shocks $\varepsilon.$, such that: (i) equation (A.1) holds for any alternative choice of task sets $(\mathcal{T}_l', \mathcal{T}_m')$; (ii) equation (6) holds at any point in the distribution (that is, firms make zero profits); (iii) the marginal distribution of occupations conditional on wages follows equation (8) (that is, workers optimize); and (iv) the unconditional marginal distributions of skills s and occupational shocks ε follow $\mathcal{N}(\bar{s}, \Sigma_s)$ and $\mathcal{N}(0, \varsigma^2 I)$, respectively.

For our quantitative exercises, we assume that the productivity of any automated task $\tau^\star$ equals $\bar{z}_{\tau^\star}$.

A.2 Deriving the Laplace approximation of the skill posterior

In order to obtain the expressions presented in the main text, it is necessary to compute the posterior skill distribution of workers who choose occupation o in a given period. Because the exact distribution is not tractable, we rely on Laplace approximations. Laplace approximations provide a normal approximation of the true posterior. To apply this method, we write down the exact un-normalized posterior log density of workers in occupation o:

$$\phi_o(s) = -\frac{1}{2}(s-\bar{s})'\Sigma_s^{-1}(s-\bar{s}) + \nu^{-1}\mu_o + \nu^{-1}A_o \cdot s - \log\left(\sum_{o'}\exp(\nu^{-1}\mu_{o'} + \nu^{-1}A_{o'}\cdot s)\right).$$

The Laplace approximation uses the posterior mode as the mean of a multivariate normal and the score of the posterior likelihood as its co-variance matrix. Thus, we need to find the first and second derivative of the un-normalized posterior. Defining the skill-conditional employment share as $h_o(s) = \frac{\exp(\nu^{-1}\mu_{o'}+\nu^{-1}A_{o'}\cdot s)}{\sum_{o''}\exp(\nu^{-1}\mu_{o''}+\nu^{-1}A_{o''}\cdot s)}$, we can write:

$$\nabla\phi_o(s) = -\Sigma_s^{-1}(s-\bar{s}) + \nu^{-1}A_o^\intercal - \nu^{-1}\sum_{o'}h_{o'}(s)A_{o'}^\intercal$$

$$\nabla^2\phi_o(s) = -\Sigma_s^{-1} - \nu^{-2}\sum_{o'}h_{o'}(s)A_{o'}^\intercal A_{o'} + \nu^{-2}\left(\sum_{o'}h_{o'}(s)A_{o'}^\intercal\right)\left(\sum_{o'}h_{o'}(s)A_{o'}^\intercal\right)^\intercal$$

The approximate posterior mean sets the first derivative to zero:

$$\bar{s}_{|o} = \bar{s} + \nu^{-1}\Sigma \overbrace{\left(A_o^\intercal - \sum_{o'}h_{o'}(\bar{s}_{|o})A_{o'}^\intercal\right)}^{\text{relative task intensity of occupation } o},$$

which implicitly defines $\bar{s}_{|o}$.
The posterior covariance matrix is then

$$\Sigma_{s|o} = -\nabla^2\phi_o(\bar{s}_{|o})^{-1} = \left(\Sigma_s^{-1} + \nu^{-2}\underbrace{\left(\sum_{o'}h_{o'}(\bar{s}_{|o})A_{o'}^\intercal A_{o'} - \left(\sum_{o'}h_{o'}(\bar{s}_{|o})A_{o'}^\intercal\right)\left(\sum_{o'}h_{o'}(\bar{s}_{|o})A_{o'}^\intercal\right)^\intercal\right)}_{\text{Task intensity dispersion across occupations}}\right)^{-1}.$$

A.3 Identification proof

In this section, we prove the following claim:

Proposition A.1. *The model is identified in steady state – that is, for two distinct parameter tuples* $\theta = (\bar{s}, \Sigma_s, \nu, \varsigma) \neq (\bar{s}', \Sigma_s', \nu', \varsigma') = \theta'$*, the model-implied distributions* F, F' *of observable data are distinct and thus the parameters can be distinguished with infinite data.*

For the proof, we assume that A has full rank, which is a condition that is satisfied with the matrix A we consider in the quantitative exercises.

Proof: By way of contraposition, suppose that the distribution of observables is identical for two models, so $F = F'$. We will show that this implies $(\bar{s}, \Sigma_s, \nu, \varsigma) = (\bar{s}', \Sigma_s', \nu', \varsigma')$. We begin by considering an arbitrary observed worker. Note that the worker is observed infinitely many times in every occupation, so $\mathbb{E}[w_o|s]$ is observable for all o and thus $\mathbb{E}[w_\cdot|s] - \mu_\cdot = As$ is identified. Since A is full rank, s is identified for that worker. Since there are infinitely many such workers, $\bar{s}, \Sigma_s$ are identified from F because the multivariate normal distribution is well-identified. Since $F = F'$, it thus follows that $\bar{s} = \bar{s}'$ and $\Sigma_s = \Sigma_s'$.
Next, for an arbitrary pair of occupations o, o' with distinct loadings $A_o \neq A_{o'}$, it is always possible to find a worker with skill s such that $(A_o - A_{o'})s \neq 0$ and therefore the worker's relative log choice probability $\log P(\hat{o} = o) - \log P(\hat{o} = o') = \nu^{-1}(A_o - A_{o'})s$ identifies ν uniquely from F. Thus, $\nu = \nu'$.
Lastly, the variance of wages over time for an arbitrary worker with skill s in occupation o is given by ς^2, which is therefore identified from F. Therefore, $\varsigma = \varsigma'$, concluding the proof. ■

B Empirical appendix

B.1 Historical examples of job transformation

This section illustrates job transformation using two historical case studies. First, consider weavers' work over the course of the 19th century. We draw on Bessen's (2011) analysis of records from the Lawrence Company, which operated several mills in Massachusetts. Table B.1 shows how weavers' duties shifted away from a diverse set of manual tasks like preparing, dressing, and letting off the reed, toward spending more time on a narrower set of tasks, like fixing mechanical issues or replacing inputs such as bobbins.

Second, machinists in manufacturing experienced job transformation during the late 20th century. Bartel *et al.* (2007) study how investments in IT-enhanced machinery transformed skill requirements in valve manufacturing. These enhancements consisted of capital-embedded

Period	Preparatory tasks		Tasks while machine running							Tasks while power loom stopped							
	Prepare warp	Dress warp	Let off warp	Pick shuttle	Beat reed	Take up cloth	Adjust warp tension	Replace empty bobbin	Monitoring	Fix smashes	Adjust temples	Back up loom	Replace empty shuttle	Fix broken weft	Fix broken warp end	Remove cloth, cleaning	Replace warp
Handloom	●	●	●	●	●	●	●		●		●		●	●	●	●	●
Early power loom (~1820)							●	●	●	●	●	●	●	●	●	●	●
1833							●	●	●	●		●	●	●	●	●	●
1883							○	●	●	●			●	●	●	●	○

Table B.1: Job transformation: the case of weavers

Notes. Legend: ● = Task performed; ○ = Reduced frequency; Empty = Task not performed.
Replication of Bessen (2011, Table 2).

improvements in computer numerically controlled (CNC) tools, which fix the raw material and automatically machine valve components based on designs entered into the machine's operating software. Previously, operators used routine machining skills and physical strength to position valves correctly and had to choose and move appropriate cutting tools. With more sophisticated CNC machines, work shifted toward setup, monitoring, and adjustment – tasks demanding greater programming and problem-solving skills. Atalay *et al.*'s (2020) systematic analysis of job advertisements corroborates this account. They show that machinists experienced major task shifts in the 1980s with CNC adoption, with increased emphasis on computer programming, problem-solving, and technical engineering knowledge.

B.2 Clustering of occupation-specific tasks

This section describes in more detail how we construct the task clusters on which our analysis is based. Our starting point is the list of detailed task statements in the O*NET database. We retain only the subset of tasks that O*NET classifies as "core" for any given occupation, dropping "supplemental" tasks that are less relevant and/or important to the occupation.

B.2.1 Extracting skill requirements

For each of the detailed tasks, we extract a core activity and skill requirements using a LLM, specifically openAI's GPT-4o model. The system and user prompts are stated below.[B.1]

Table B.2 provides an illustration of the core activity and required skills this approach extracts for a set of skills; it also indicates the cluster the cluster will eventually be assigned to.

System Prompt.

```
<Role> You are an expert in labor economics, job analysis, and task
   classification. </Role>

<Overall goal>
You will be presented with a list of {len(tasks_chunk)}
   occupation-specific task statements. The ultimate goal is to group
   these and thousands of other tasks into clusters based on the type
   of activity and skills utilized, i.e. someone skilled at one task
   in a cluster could perform others in that cluster.

Your overall task is to prepare this clustering step by identifying,
   for each task statement:
(i) the fundamental  work activity; and (ii) the most essential skills
   and abilities (up to 5) required to perform this task effectively.

Key requirements for (i) the fundamental work activity:
- Definition: The fundamental work activity is a concise, abstract
   description that encapsulates the core activity  involved in the
   task statement (what is being done).
- Generalization: The activity label should be broad enough that if
   someone can perform one task under this label,
```

[B.1] The *temperature* parameter is set to 0.00001, which directs the model to provide its most confident response, minimizing variation across runs. Technically, the LLM predicts the next word in a sequence based on the preceding words and its prior training. Denoting by q_i the logit for candidate token i, the softmax function is used to scale the logits and map them into probabilities: $\frac{e^{q_i/T}}{\sum_{k=1}^{\text{size of vocabulary}} e^{q_k/T}}$. The parameter T is known as *temperature*. A higher temperature value "excites" previously low probability candidates, encouraging creativity, whereas a lower temperature value lowers the smaller outputs relative to the largest outputs. A lower value is thus preferable for contexts requiring high coherence and accuracy. Note, though, that even $T = 0$ does not result in deterministic output in practice, likely due to sources of randomness such as the state of the random-number generator. Moreover, the so-called "reasoning" generation of models does not support a temperature parameter. In practice, we have verified that the time allocation shares are highly comparable across different runs of the model.

```
 they'd be expected to handle any task requiring that same underlying
    competency.
- Terminology: Use concise and standardized, domain-agnostic terms
    that capture the core function, phrasing them in clear,
    natural-sounding language.
- Self-explanatory: The label must offer a succinct, self-contained
    summary that includes essential context for standalone
    understanding;
      do not merely reduce the statement to a vague abbreviation.
- Predominant activity: When multiple actions are present, select the
    one that best represents the overall purpose of the task.

Key requirements for (ii) the skills and abilities:
- Definition: skills refer to developed capacities that facilitate
    performance of activities that occur across jobs; abilities refer
    to relatively enduring attributes of an individual's capability for
    performing a particular range of different tasks. A "skill" is not
    simply a rewording of the task/activity description itself, but
    rather answer the question "What underlying capability makes
    someone good at this task?"
So for each skill you identify, ask: 'Would this skill enable
    performance across MULTIPLE different tasks and contexts?' If not,
    it's likely not a true underlying skill.
- Task: Identify the essential skills and abilities required to
    perform this task effectively and list them in descending order of
    importance.
- The number of skills can range from 1 to 5, depending on the
    complexity of the task; for straightforward ones, only include the
    core skills (at least 1); avoid padding with peripheral skills.
- The "most important skills" can include both capabilities and, where
    critical, knowledge domains, including:
i) Cognitive capabilities
Examples:  strategic planning, statistical analysis, diagnostic
    reasoning, technical writing
ii) Specialized technical capabilities
Examples: programming, surgical technique, database architecture
```

```
iii) Interpersonal capabilities, management, and leadership
Examples: negotiation, leadership, instruction, conflict resolution,
   team development, performance evaluation, delegation,
   organizational design, change management,
management of financial resources, management of personnel resources
iv) Physical/sensory capabilities
Examples: fine motor control, spatial awareness, physical endurance
v) Specialized expertise areas
Examples: mathematical modeling, designing scientific experiments,
   legal precedents, medical protocols
- Each skill in the list must follow this format: "Skill Name (Level)"
- Level must be one of "basic", "intermediate", "advanced", or
   "expert" using the following criteria:
basic: requires fundamental knowledge and minimal experience;
   intermediate: requires specialized knowledge and moderate
   experience; advanced: requires deep expertise and substantial
   experience;
expert: requires mastery-level knowledge, typically 8+ years of
   focused experience.
- Critial: When identifying skills, pay particular attention to
   specialized capabilities that typically command higher wages in the
   labor market, such as: Complex analytical or strategic thinking
   skills, Specialized technical expertise that requires extensive
   training, High-stakes decision-making capabilities, Skills
   involving the direction of others' work or significant resources,
   Expertise that is both scarce and in high demand. For high-wage
   occupations, ensure you separately list these skills rather than
   using generic descriptors.
</Overall goal>

<Detailed instructions>
Step 1) For each task statement, identify and summarize (i) the
   fundamental work activity; and (ii) the most important  skills.
Step 2) Return the output (activity; skills) for all
   {len(tasks_chunk)} task statements in the JSON format specified.
- List the skills in descending order of importance to the task (most
```

```
   crucial first).
- Never leave any task blank; if unsure, provide your best guess.
</Detailed instructions>

<Examples>
The following examples illustrate the level of abstraction desired
   (for reference only, do not copy these exact labels unless they
   truly match the task at hand).

Example work activities: "train and teach others at work", "operate
   vehicles", "operate industrial machinery," "provide advice or
   consultation," "coordinate the work of subordinates/peers,"
   "inspect or repair equipment," etc.

Example task: "Review statistical studies, technological advances, or
   regulatory standards and trends to stay abreast of issues in the
   field of quality control."
Activity: evaluate complex technical information
Skills: analytical thinking (expert), research (advanced), statistical
   analysis (advanced), reading comprehension (advanced)

Example task: "Wash glasses or other serving equipment at bars."
- Activity: cleaning
- Skills: manual dexterity (basic)

Example task: "Analyze financial statements to determine company
   valuation"
Activity: analyze and interpret financial data
Skills: market analysis (expert), numerical reasoning (advanced), data
   analysis (intermediate), financial modeling (advanced)

Example task: "Lead strategic planning for a multinational division
   with $500M annual revenue"
Activity: direct organizational strategy
Skills: leadership (expert), strategic planning (expert), financial
   analysis (advanced), business intelligence (advanced)
```

```
Example task: "Train new employees on safety procedures and equipment
   operation"
Activity: train and teach colleagues
Skills: verbal communication (intermediate), technical knowledge about
   equipment (intermediate), instructional planning (basic)

Example task: "Supervise and coordinate the work plan of customer
   service representatives and schedule shifts"
Activity: manage team operations
Skills: operational planning (intermediate), verbal communication
   (advanced), people development (advanced)

Example task: "Develop marketing strategy for new product launch"
Activity: create marketing strategies
Skills: strategic thinking (advanced), business knowledge (expert),
   creativity (intermediate), analytical reasoning (intermediate),
   written communication (advanced)

Example task: "Read operating schedules or instructions or receive
   verbal orders to determine amounts to be pumped."
- Activity: follow operational instructions
- Skills: reading comprehension (basic), verbal communication (basic)
</Examples>
```

User Prompt.

```
<List of task statements>
Here is the list of {len(tasks_chunk)} job task statements to analyze,
   along with their index numbers:

{task_list}
</List of task statements>
```

Task	Activity	Skills	Cluster
Smooth rough spots on walls and ceilings, using sandpaper	smooth surfaces	manual dexterity (basic), attention to detail (basic)	Performing Detailed Manual Tasks
Lubricate moving parts on gate-crossing mechanisms and swinging signals	lubricate moving parts	manual dexterity (basic), attention to detail (basic)	Performing Detailed Manual Tasks
Perform physically demanding tasks, such as digging trenches to lay conduit or moving or lifting heavy objects	perform physical labor	physical endurance (advanced), manual dexterity (intermediate)	Performing Physical Labor
Prepare reports of activities, evaluations, recommendations, or decisions	prepare reports	report writing (advanced), analytical reasoning (intermediate), attention to detail (intermediate)	Processing and Analyzing Records
Confer with officials of public health and law enforcement agencies to coordinate interdepartmental activities.	coordinate interdepartmental activities	collaboration (advanced), project management (advanced), communication skills (intermediate)	Coordinating Project Initiatives

Table B.2: Examples: detailed tasks, extracted characteristics, and cluster assignment

Notes. This table lists examples of detailed tasks (first column), that is the input, as well as the extracted core activity and skill requirements (LLM-generated), and the labeled cluster to which this task is assigned.

B.2.2 Embeddings and clustering

We use Alibaba's `gte-Qwen2-1.5B-instruct` model to create word embeddings of dimension 1,536 for the extracted skills for each task statement. To prepare the embeddings data for clustering, and noting that the HDBSCAN algorithm we are using performs best on data with low to medium dimensionality, we next perform a two-part dimensionality reduction step. We initially perform PCA, retaining the principal components that explain 95% of the variance in the embedding space. We then perform a subsequent UMAP step, which is useful to preserve both local and global data structures while shrinking the number of components to a level suited for the HDBSCAN algorithm. Finally, we use the HDBSCAN algorithms with the following hyperparameters `min_cluster_size` = 70, `min_samples` = 40, `cluster_selection_epsilon` = 0.05. The distance metric er `option_hdbscan_metric` is Euclidean given the preceding UMAP step.

B.2.3 Labeling step & summary output

Finally, we use OpenAI's o3-mini-high model to create natural-language labels and a summary description for each of the task clusters. These cluster-level meta data are useful in two ways: in terms of interpretation, and as inputs to the LLM when constructing the occupation-level time shares across the task clusters. Practically, for each cluster we randomly select ten representative tasks and feed the core activity as well as the skill requirements for these tasks to the LLM, instructing it to generate a cluster label and a brief description, per the following prompts.

Table B.3 details all 38 task clusters, indicating the summary label and description.

System Prompt.

```
 <Role> You are a world-class expert in labor economics, task
    classification and  occupational analysis.  You use concise and
    standardized language that is consistent with established
    terminology in skills/occupational databases like O*NET or PIACC.
</Role>

<Overall goal>
The overarching goal is to create accurate and meaningful summary
    labels for clusters of job tasks.
Each cluster comprises many tasks, which grouped by the type of
    activity (what is being done) and the skills required (capacities
    that facilitate performance of activities); i.e., the general rule
    is that a person proficient in one task in a given cluster should
    also be able to perform others in that cluster.

Given this goal, you will be presented with a list of tasks --
    alongside the most essential skills required to perform each --
    that exemplify a particular cluster.

You primary task is to create an accurate and concise summary label
    for this cluster of tasks.
Your secondary task is to provide a concise description of this
    cluster, with reference to core skill requirements differentiating
    this cluster from others.

Requirements for the summary label:
```

```
- The label summarizes the common core activities (what is being
   done), while remaining specific enough to meaningfully
   differentiate this cluster from others.
- The label focuses on the essential underlying activity rather than
   the specific domain.
- The label is sufficiently specific to allow differentiating between
   occupations that have different skill requirements and wage levels.
- The label is concise (2-5 words),  uses natural sounding language
   aligned with established task/skill terminology,
    and where possible begins with a gerund (verb+ing form).

Requirements for the description:
- The concise description (1 sentence) summarizes the cluster, with
   reference to core skill requirements differentiating this cluster
   from others.
</Overall goal>

<Detailed instructions>
Step 1: Analyze the {len(tasks_chunk)} tasks by identifying the
   fundamental activities involved and core skills utilized across all
   them.

Step 2: Create a summary cluster label that satisfies the requirements
   outlined above.
Test your label to ensure that it meets each of the X requirements;
   revise and iterate until this is the case.

Step 3: Given the label, and considering the skills listed for the
   exemplary rasks, provide a concise description.
</Detailed instructions>

<Examples of cluster labels>
Here are examples of cluster labels to illustrate the desired level of
   abstraction. These serve for guidance only, you must create
   appropriate task-specific labels.
```

```
- Positive example: "Developing and Building Teams" (relevant across
   domains, but specific enough to distinguish from other
   interpersonal tasks)
- Positive example: "Analyzing quantitative data" (relevant across
   different occupations, distinct from qualitative analysis which
   would involve different skills)
- Positive example: "Performing gross motor or heavy manual physical
   labor" (connotes a broad range of tasks with similar skill
   requirements)
- Positive example: "Technical Operation and Maintenance Tasks" (not
   domain specific, connotes a skill requirement distinct from
   advanced technical analysis)
- Negative example: "Getting Information" (too unspecific)
- Negative example: "Performing Administrative Activities" (too broad,
   could involve routine tasks such as processing paperwork or
   advanced managerial tasks, i.e. tasks requiring very different
   skills)
- Negative example: "Communication" (too unspecific, could comprise
   anything from chatting with colleagues to arguing a complex case in
   court)
</Examples>
```

User Prompt.

```
 <List of tasks to analyze>
Here is the list of {len(task_descriptions)} tasks that are
   representative of the task cluster under consideration alongside
   the most important skills required to perform them:

{task_list}
</List of tasks to analyze>
```

B.3 LLM-generated time diaries

This section describes how, given the task clusters, we construct the occupational task weight matrix. In addition, we detail validation exercises.

Cluster label	Description
Performing Detailed Manual Tasks	This cluster involves executing precise, hands-on operations—ranging from cleaning and lubricating to marking and packaging—that rely on basic manual dexterity and careful attention to detail.
Performing Precision Finishing Tasks	This cluster encompasses tasks that involve fine manual adjustments and finishing operations—such as aligning, smoothing, and testing components—requiring intermediate manual dexterity and attention to detail.
Preparing and Planning Meals	This cluster involves tasks that span cooking, menu planning, and overseeing food safety and service, requiring strong culinary skills, dietary knowledge, and attention to detail.
Maintaining Records and Inventory	This cluster involves routine operational support tasks that require diligent record keeping, inventory management, and clear communication to sustain documentation, asset tracking, and service functions.
Coordinating Detail-Oriented Operations	This cluster involves routine tasks such as sorting, record-keeping, material distribution, and facility upkeep that require meticulous attention to detail and basic to intermediate organizational skills.
Delivering Public Presentations	This cluster involves speaking in formal and public settings—ranging from project briefings and lectures to courtroom testimonies—requiring advanced public speaking, communication, and subject matter expertise.
Documenting Technical Information	This cluster focuses on capturing and recording technical details and processes using advanced technical writing, documentation, and attention to detail.
Performing Clinical Procedures	This cluster involves executing patient-focused clinical tasks that combine advanced diagnostic reasoning, technical equipment operation, interpersonal communication, and therapeutic interventions to assess and treat medical conditions.
Providing Customer Service	This cluster involves direct customer interactions that require strong interpersonal, communication, time management, and organizational skills to assist, guide, and support various client needs in service-oriented settings.
Administering Regulatory Compliance	This cluster involves interpreting policies, reviewing and enforcing regulatory standards, and developing procedures, all requiring advanced regulatory knowledge, analytical reasoning, and communication skills.
Coordinating Emergency Response	This cluster involves executing and managing emergency procedures, crisis communication, threat monitoring, and strategic planning, requiring advanced emergency response and situational awareness skills.
Maintaining and Managing Records	This cluster involves systematically updating, retrieving, and organizing diverse records and data through strong attention to detail and organizational skills.
Reviewing and Editing Information	This cluster involves accurately reviewing, editing, and verifying various forms of information—from written materials to operational data—requiring advanced attention to detail and precision.
Ensuring Regulatory Compliance	This cluster involves meticulous inspection, record management, and analytical review to verify adherence to regulatory standards and operational protocols.
Performing Physical Labor	This cluster encompasses physically demanding tasks that require manual dexterity, physical endurance, and fundamental technical and safety skills across diverse settings including construction, cleaning, material handling, animal care, and exercise instruction.
Creating Technical Visual Representations	This cluster involves transforming data, technical specifications, and artistic ideas into precise visual media by integrating advanced drafting, design, and multimedia editing skills.
Designing and Implementing Systems	This cluster centers on planning, designing, and integrating technical systems across diverse fields, emphasizing advanced project management, engineering design, and technical expertise.
Processing and Analyzing Records	This cluster involves tasks focused on maintaining, recording, and evaluating data—including financial, production, and medical records—where strong numerical reasoning, analytical skills, and meticulous attention to detail are essential.
Operating, Calibrating, and Inspecting Equipment	This task cluster involves technical operations focused on handling electronic recording, imaging, and sound equipment, requiring precise calibration, systematic inspections, and adept problem-solving skills.
Inspecting and Evaluating Quality	This cluster involves detailed inspections and analyses that rely on advanced analytical reasoning and attention to detail to assess product, site, and process quality, ensuring standards and performance are met.
Performing Skilled Manual Operations	This cluster involves executing diverse manual tasks—ranging from assembly, finishing, and equipment maintenance to operation and cleaning—that require intermediate to advanced manual dexterity, attention to detail, and technical proficiency.
Negotiating and Coordinating Contracts	This cluster involves engaging stakeholders through advanced negotiation and communication skills to secure agreements and manage procurement activities while coordinating legal, regulatory, and project management requirements.
Repairing and Maintaining Equipment	This cluster encompasses preventative maintenance, technical repair, and equipment installation tasks that require advanced system knowledge, manual dexterity, and safety awareness.
Managing Safety Operations	This cluster involves overseeing operational activities with a strong emphasis on safety compliance, hazard assessment, and technical oversight across diverse industrial, emergency, and technical settings.
Monitoring and Inspecting Systems	This cluster involves actively operating, adjusting, and inspecting automated processes and equipment by employing advanced technical troubleshooting, precision measurement, and quality control skills to ensure optimal system performance.
Analyzing and Optimizing Systems	This cluster involves applying advanced technical analysis, simulation, and maintenance skills to assess performance, recommend design changes, and ensure operational integrity across diverse systems.
Analyzing Natural Phenomena	This cluster involves applying advanced scientific analysis, technical expertise, and data interpretation to evaluate, classify, and redesign natural and biological systems across diverse domains.
Instructing and Training	This cluster involves delivering instruction, training, and mentorship across diverse subject areas, relying on advanced instructional techniques, verbal communication, and subject matter expertise.
Mediating and Consulting Clients	This cluster involves interpersonal guidance tasks—including counseling, referrals, conflict investigation, and dispute resolution—that require advanced communication, empathy, and problem-solving skills to address diverse client issues effectively.
Developing and Delivering Instruction	This cluster encompasses tasks centered on planning, designing, and conveying educational programs and curricula, leveraging advanced instructional design, curriculum development, and communication skills across varied content areas.
Communicating and Educating	This cluster involves effectively conveying information, instructions, and feedback through verbal channels, integrating clear reporting, problem-solving, and instructional skills across diverse contexts.
Engaging in Continuous Learning	This cluster encompasses tasks that require ongoing research, information synthesis, and professional development to remain current with industry trends, technology advancements, and scientific progress.
Collaborating Across Functions	This cluster comprises tasks requiring effective teamwork, communication, and coordination across diverse professional areas to address problems, manage operations, and support technical and client-oriented activities.
Coordinating Project Initiatives	This cluster involves planning, overseeing, and collaborating on diverse project tasks, leveraging advanced project management, communication, and leadership skills.
Coordinating Administrative Tasks	This cluster encompasses planning, scheduling, and organizing a range of administrative operations, requiring strong organizational, communication, and project management skills.
Coordinating Strategic Initiatives	This cluster involves planning, organizing, and supervising diverse activities—ranging from educational events to disaster recovery and recruitment—requiring advanced leadership, strategic planning, and team management skills.
Producing Technical Documentation	This cluster involves drafting and compiling technical reports, proposals, and documentation through advanced technical writing, analytical reasoning, and data presentation skills, with elements of programming and research support.
Performing Strategic Analysis	This cluster involves advanced quantitative research, financial and cost analyses, and strategic planning to assess deviations, forecast outcomes, and drive management recommendations.

Table B.3: Task cluster labels and descriptions

B.3.1 Methodology

To generate the time diaries we use the latest version of GPT-o3-mini-high for this step. We loop over each occupation using the following prompts. These are designed to break the complex task into clear sequential steps, draw on high-quality inputs, and convert the qualitative assessment into a numerical output.

System Prompt.

```
You are an expert in occupational classification (the system being
   used is {occ_system} at the {occ_level}-digit level) and analyzing
   occupational time allocation.
You combine precision in classification with deep knowledge of how
   different occupational groups allocate their time across tasks.
You focus on accurate, structured data output, and your time share
   predictions MUST sum to exactly 1.0. You are precise and
   conscientious.
```

User Prompt.

```
<Objective and context>
We want to accurately estimate what percentage of their work time
   workers in a specific occupation group spend on various tasks.
The occupation group is {occ_title}, as classified following the
   occupational classification system {classification_description}.
The reference period to consider is the {option_timeperiod}.
</Objective and context>

<List of tasks>
The tasks to consider are as follows:

{task_list}
</List of tasks>

<Instructions>
Follow these steps to generate accurate time allocation estimates:

1. Analyze core functions, activities and responsibilities of
   {occ_title}
```

```
2. For each task listed above:
  - Review the task carefully
  - Assess the importance and frequency of this task for {occ_title}
   in the {option_timeperiod}, drawing on high-quality evidence,
   expert knowledge and statistical data.
3. Having done this for all tasks, convert assessments to time
   allocation shares:
    For each task:
  - Convert assessment to percentage of work time
  - Translate to decimal (e.g., 25%  0.2500)
  - Document: task_name: 0.XXXX
  - Add to running_sum
4. Verification (required):
  Calculate total_sum to 4 decimals
  If total_sum != 1.0000:
  - Calculate scaling = 1.0000/total_sum
  - Multiply EACH share by scaling
  - Recalculate sum
  STOP: Submit shares only if sum = 1.0000

Critical Requirements:
- Use 4 decimal precision throughout
- Show calculations
- Final shares MUST sum to 1.0000
- No rounding of intermediate values
- Calculate time shares for all {task_count} tasks.
</Instructions>
```

In rare instances, the LLM does not generate time shares that sum to 1, despite the above instructions. This is reminiscent of human responses in time diary surveys. We therefore programmatically normalize the LLM-predicted shares, just as we do using the conventional, human survey responses discussed below. Over the course of the project we moved from openAI 4o to o3-mini-high, which dramatically reduced the need for this ex-post normalization of time shares.

B.3.2 Validation

This section describes three complementary approaches that collectively demonstrate the robustness of the LLM-based measurement of the occupational task weights. Additionally, Appendix C.2.4 describes in greater detail an alternative that relies on aggregated O*NET task importance weights.

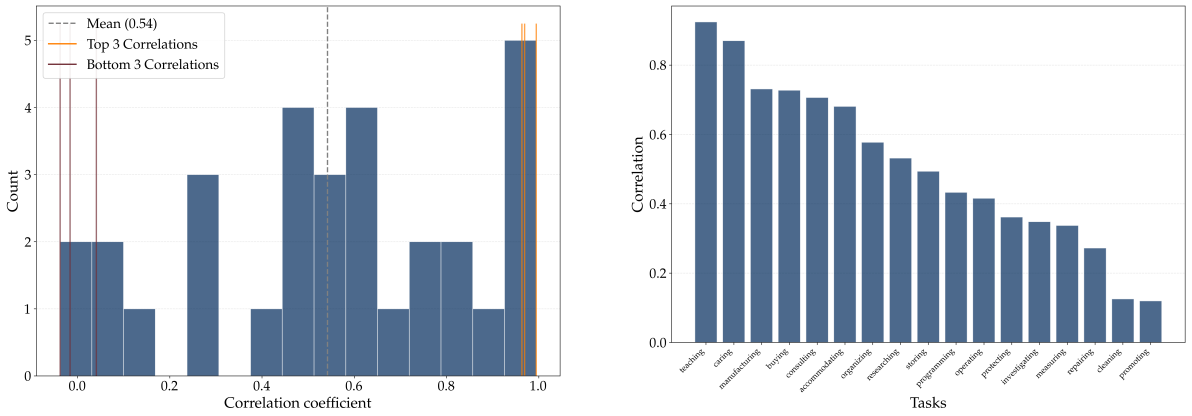
BIBB time diaries. We use a supplemental survey conducted for the 2012 Employment Survey carried out by the German Federal Institute for Vocational Training (Bundesinstitut fuer Berufsbildung, BIBB) and the and the German Federal Institute for Occupational Safety and Health (BAuA). This survey asks a subset of surveyed workers to report their allocation of time to a pre-specified list of tasks such as "teaching" and "cleaning" on a given day.[B.2]

We proceed in three steps. First, we construct occupation-task level time allocation shares from the BIBB. We consider the sample of individuals in West Germany aged 15-65 who have completed their training and who report a valid occupation ISCO-08 2-digit occupation. For each individual, we normalize the time shares to sum to one. Then we average time shares across occupations and retain those occupations comprising least ten surveyed workers. Second, we re-run the same LLM-based process as in the main analysis, but now requiring responses for the same set of tasks considered in the BIBB and looping over ISCO-08 2-digit occupations. Third, we compare the BIBB-based and LLM-based responses.

Overall, the two different approaches yield highly comparable results. Figure B.1a shows correlations at the level of occupations; the mean correlation is 0.54, the standard deviation is 0.31. The lowest correlations are reported for "Customer Service Clerks," "General and Keyboard Clerks" and "Numerical and Material Recording Clerks." A major source of discrepancy is that for these occupations, survey respondents in the BIBB put substantial weight on the task "programming" (which in the original German language context could also be interpreted as "using a computer"). With further clarification on the interpretation of tasks, we expect that the LLM and BIBB would yield results that are more comparable still. Figure B.1b reports task-level correlations of weights across occupations. The tasks with the lowest overlap are "promoting" and "cleaning," while the alignment is greatest for "teaching" and "caring."

Comparison to O*NET importance weights for GWAs. Next, we compare LLM-generated time allocation shares with O*NET occupation-level importance weights for Generalized Work Activities (GWAs). We use the GWAs from ONET 5.0, as this database aligns with the SOC-2000 classification used in our main empirical analysis. We construct relative importance weights for each GWA by occupation and aggregate to the minor-group level. We then generate LLM-based

[B.2] The full list of 17 tasks is as follows: 'investigating', 'organizing', 'researching', 'programming', 'teaching', 'consulting', 'buying', 'promoting', 'repairing', 'accommodating', 'caring', 'cleaning', 'protecting', 'measuring', 'operating', 'manufacturing', 'storing'



(a) Occupation-level correlations

(b) Task-level correlations

Figure B.1: Comparison of LLM-generated task weights & BIBB survey

Notes. The left panel plots the distribution of occupation-level correlations between the LLM-predicted task weights and those constructed from the BIBB. The right panel shows the task-level correlations across occupations instead.

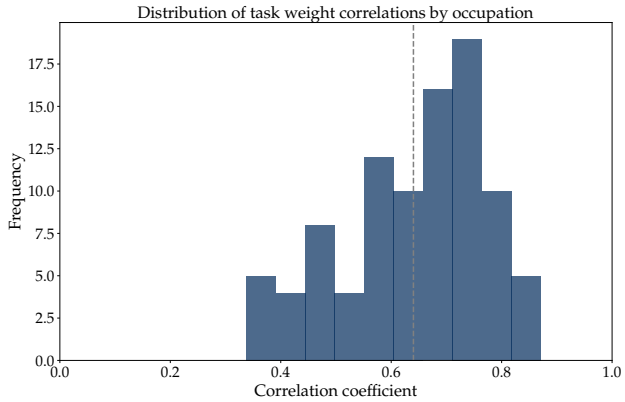


Figure B.2: LLM-time shares for GWAs correlate with O*NET importance weights

time allocation shares for identical GWAs across the same occupational categories and compare the resulting two A matrices.

Figure B.2 displays the distribution of occupation-level correlations between LLM-generated time shares and O*NET importance weights. The distribution is markedly right-skewed, with a central tendency around 0.6-0.7, indicating substantial alignment between our LLM-based approach and established occupational measurements. Figure B.3 presents task-specific correlations across occupations, grouped by correlation strength. Tasks involving cognitive and managerial functions show the strongest correspondence (correlations >0.75), while more specialized technical tasks exhibit moderate alignment. Even the lowest-correlating tasks maintain coefficients above 0.2, suggesting our approach captures meaningful variation across the entire task spectrum.

LLM consistency in aggregation across occupational hierarchies. Figure B.4 demonstrates that

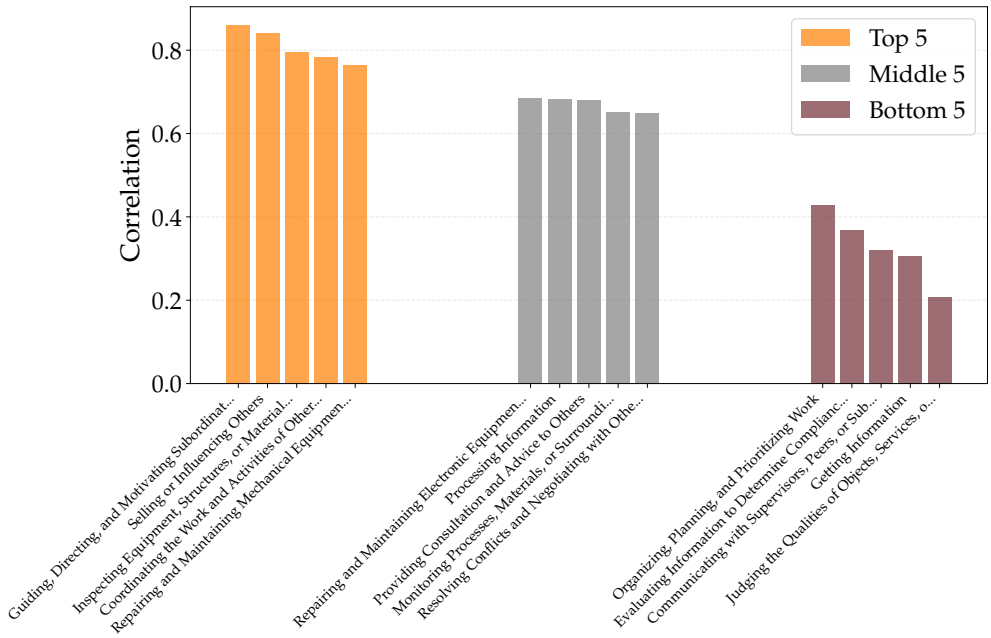


Figure B.3: Correlation across occupations by task

the LLM-generated task weights are consistent across different levels of occupational aggregation. We compare weights derived directly from major occupational groups with those constructed by averaging across their constituent minor occupational categories (using unweighted means). The very high correlation coefficient (0.89) confirms that our task approach maintains consistency regardless of aggregation level.

B.4 Occupational labor shares

To construct occupation-level labor shares – i.e., compensation over value-added – we take the following approach, where industries are indexed by j and occupations by o:

(i) Construct weights s_{oj} corresponding to the share of industry-j payments to labor going to occupation o:

$$s_{oj} = \frac{\text{(wage payments to } o \text{ in } j)}{\sum_o \text{(wage payments to } o \text{ in } j)} \tag{B.1}$$

(ii) Assume that value-added in j due to o is proportional to s_{oj}:

$$\text{VA}_{oj} = s_{oj} \cdot \text{VA}_j \tag{B.2}$$

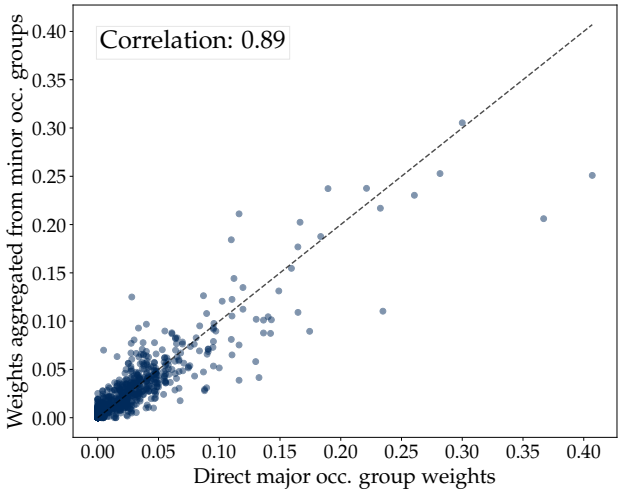


Figure B.4: Comparison of task weights at different occupational levels of aggregation

(iii) Compute

$$\text{LS}_o = \frac{\sum_j \text{wage payments to } o \text{ in } j}{\sum_j \text{VA}_{oj}} \tag{B.3}$$

In the empirical implementation, we use the 2002 wave of the BLS Occupational Employment and Wage Statistics (OEWS) to construct s_{oj}. This wave uses the same SOC-2000 occupational classification as our (harmonized) NLSY dataset and NAICS-2002 industry codes. Data on VA_j come from the BEA-BLS Integrated Industry-level Production Accounts (1987-2020). In addition, to construct the numerator of equation (B.3) we use the same apportionment method as in equation (B.2).[B.3] Industry-level data are averaged across sample years. We then link the OEWS data on s_{oj} with the BEA/BLS industry-level data by merging at the 2-digit NAICS level, retaining only those industries with a 1:1 mapping.[B.4]

The (unweighted) average labor share across occupations is 0.61, with a minimum of 0.49 (Farming, Fishing, and Forestry Occupations) and a maximum of 0.75 (Legal Occupations).

B.5 Automation exposure measures: Webb (2019)

Figure B.5 shows the average standardized exposure score of each task cluster for the three types of technology considered by Webb (2019): AI, Robots, and Software.

It can be observed that the task cluster identified as most exposed to AI in Webb (2019) is

[B.3] Using the wage bill information from the OEWS instead suffers from the problem that magnitudes of compensation differ from those in the BEA-BLS accounts; using the latter is, therefore, internally more consistent.

[B.4] The BEA/BLS data provide a crosswalk from the "production account classes" to NAICS-2007; NAICS-2007 and NAICS-2002 are identical at the 2-digit level.

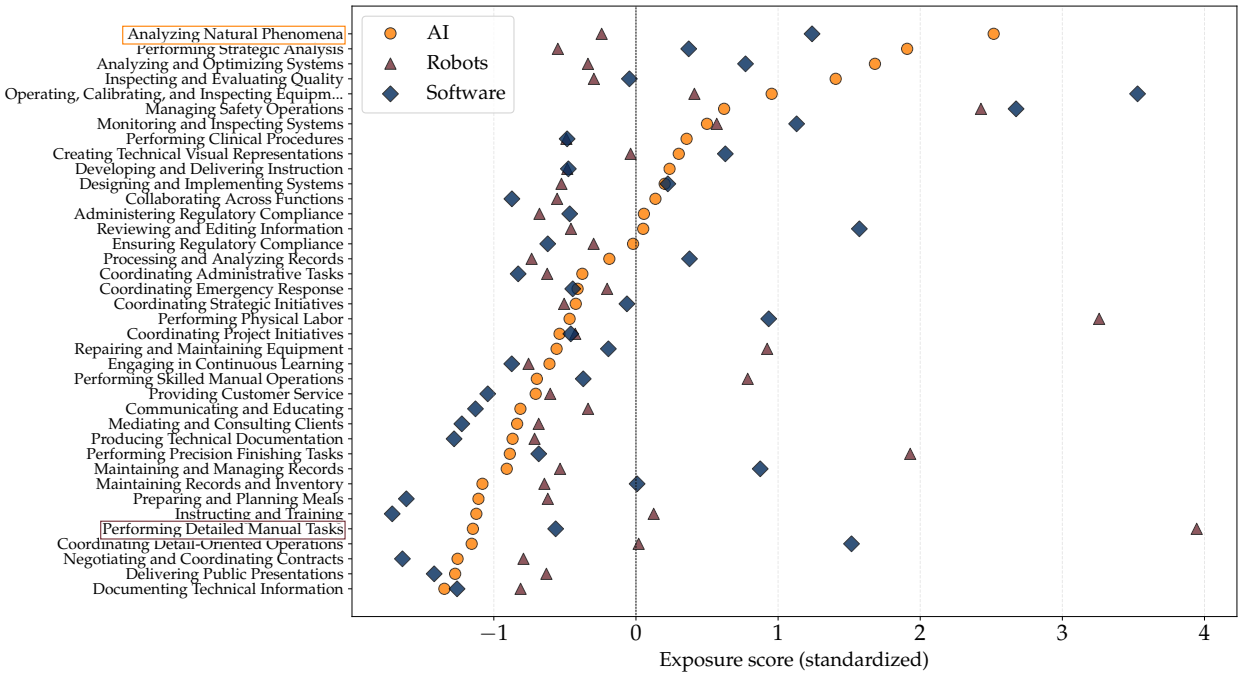


Figure B.5: Technology-specific exposure scores at the task level (Webb, 2019)

Notes. This chart shows for each of three technologies the standardized exposure score for our task clusters based on the results of Webb (2019). In Webb's approach, AI patents are identified by terms like "neural network," "deep learning," or "generative model" in titles or abstracts. Software patents contain terms such as "software" or "program" while excluding hardware-related terms like "chip" or "circuit." Robot patents are selected through the inclusion of the term "robot" in titles or abstracts.

"Analyzing Natural Phenomena" which, according to the LLM's summary description, involves "applying advanced scientific analysis, technical expertise, and data interpretation to evaluate, classify, and redesign natural and biological systems across diverse domains." By contrast, "Processing and Analyzing Record," our primary example of a task category exposed to LLMs has a close to average exposure score.

This difference is indeed to be expected. The technology cluster labeled by Webb (2019) as "AI" comprises a broader set of tools, including neural networks and deep learning algorithms more broadly, compared to the study by Eloundou *et al.* (2023), which explores task-level exposure to LLMs more specifically. Thus, it is indeed to be expected that different tasks would be exposed to the two technologies, respectively. Our framework suggests that this difference is potentially important, even if the degree of automation of two distinct tasks, as measured by the decline in the labor share, for instance, may carry labor market consequences that differ in important ways depending on the bundles of tasks the exposed tasks form part of and the distribution of task-specific skills.

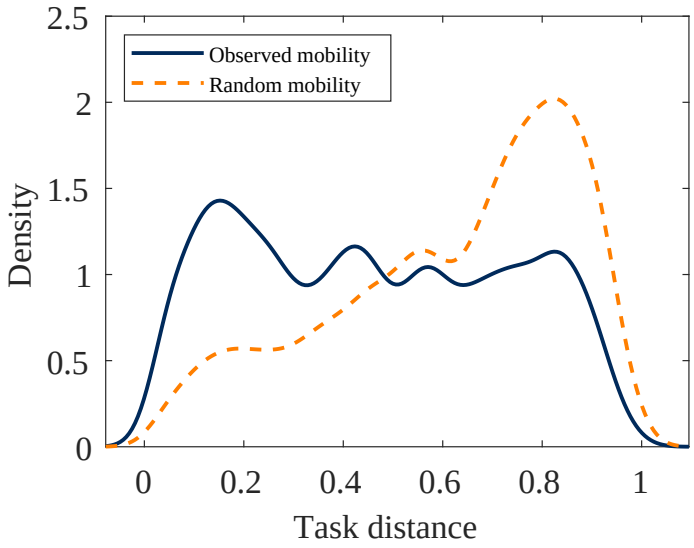


Figure B.6: Data: observed moves reflect task similarity

Notes. This figure plots the observed density of distances conditional on switching occupation (solid line) and under a random-mobility benchmark (dashed line) based on the NLSY data.

B.6 Additional results

Figure B.6 is the empirical counterpart to Figure 7b, showing that in the NLSY workers' observed switches across occupations reflect task similarity.

C Quantitative appendix

C.1 Methodology

C.1.1 Stochastic gradient descent

To estimate the model parameters, including the joint distribution of skills, the following stochastic object has to be maximized:

$$\hat{\mathcal{L}}_i(w_{i,\hat{o}_{i,\cdot},\cdot},\hat{o}_{i,\cdot}|\nu,\varsigma,\bar{s},\Sigma_s)=\left(\frac{1}{n_0}\sum_j\prod_t P(\hat{o}_{i,t}|w_{j,\cdot,t},\nu)\right)\cdot f(w_{i,\hat{o}_{i,\cdot},\cdot}|\varsigma,\bar{s},\Sigma_s)$$

To do this, we can write exploit the fact $s_i|w_{j,\hat{o}_{i,\cdot},\cdot}$ is normal and thus any can be written as

$$s_i=\mu_s^{cond}+L_s^{cond}\cdot u$$

for some easy to compute $(\mu_s^{cond}, L_s^{cond})$ and $u \sim \mathcal{N}(0, I)$. u is drawn once and then held constant throughout the maximization procedure, while μ_s^{cond} and L_s^{cond} depend on model parameters. We therefore proceed as follows:

(i) For each worker i, generate n_0 draws of u that remain fixed

(ii) Compute $s_i = \mu_s^{cond} + L_s^{cond} \cdot u$

(iii) Compute $\varepsilon_{i,t} = w_{i,\hat{o}_{i,t},t} - \mu_o - A_{o,\cdot} \cdot s_i$

(iv) Use these draws to obtain a sample $w_{j,\cdot,\cdot}$ of wages in every occupation-period cell.

(v) Compute $\hat{\mathcal{L}}_i(w_{i,\hat{o}_{i,\cdot},\cdot}, \hat{o}_{i,\cdot} | \nu, \varsigma, \bar{s}, \Sigma_s)$

We then employ stochastic gradient descent. That is, starting with some guess, we update our parameters as follows:

$$\theta_{t+1} = \theta_t - \eta \cdot \nabla \left(-\hat{\mathcal{L}}(\theta_t) \right)$$

for some sufficiently small $\eta > 0$. To further ease the computational load, we evaluate the likelihood at a subsample B only. To do this, we iterate over epochs. In each epoch, we randomly partition individuals into n groups:

$$\{1, 2, \ldots, I\} = B_1 \cup B_2 \cup \ldots \cup B_n, \quad B_i \cap B_j = \emptyset$$

and, on each iteration within an epoch, evaluate the likelihood based on batch $B_1, \ldots, B_n$ only. Using parallelization over individuals and auto-differentiation techniques, this reduces the computation time of the likelihood maximization procedure substantially and allows us to solve this problem even when the parameter space is very large.

C.2 Additional results

C.2.1 Additional estimation results

The co-variance matrix of our baseline estimate is plotted in Figure C.1.

C.2.2 Additional simulation results

Sorting. Figure C.2 illustrates sorting on the basis of comparative advantage, using the skill of "Coordinating Strategic Initiatives" for the occupation of "Top Executives" as an example. Since

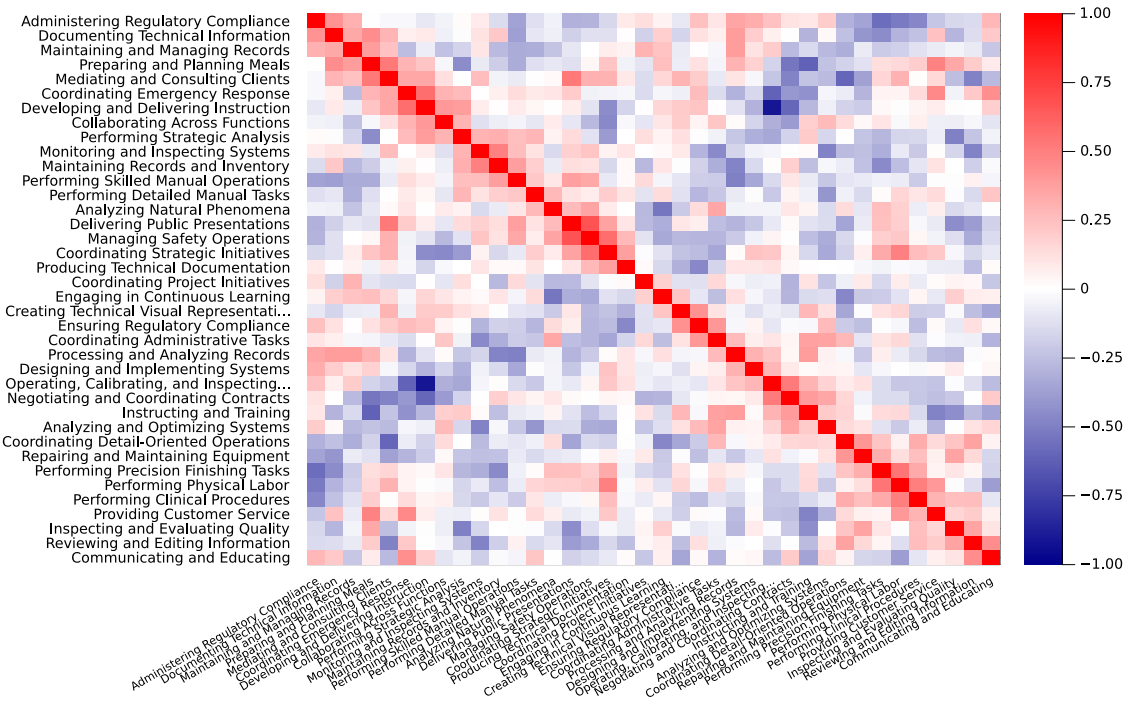


Figure C.1: Correlation matrix C_s

Notes: Heatmap of the correlation matrix C_s. Each cell corresponds to one entry of the matrix.

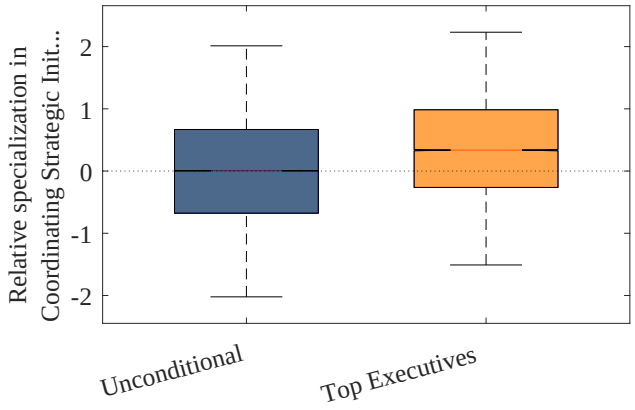


Figure C.2: Selection based on comparative advantage: example

Notes. This figure displays the standardized distribution of relative specialization in a focal task $\tau^\star$ defined as $s_{i,\tau^\star} - \frac{1}{n_{\text{skill}}} \sum_{\mathcal{T}_l} s_{i,\tau}$, where $\tau^\star$ is the task "Coordinating Strategic Initiatives," comparing the unconditional distribution that the distribution conditional on having selected into the occupation "Top Executives."

the task of "Coordinating Strategic Initiatives" is heavily utilized within this occupation, the occupation features a worker pool that is on average more skilled in this task.

Task-shift and selection effects: comparison across tasks. In the main text, we showed that the magnitude of the re-sorting effect following the automation of a particular task is positively related to the dispersion of skills in that task across workers. In this appendix, we conduct a comparative analysis of different task automation shocks to understand the determinants of task shift and selection effects. One way to assess the role of these effects with rising exposure is to regress their effect size onto the exposure $A_{o,\tau}$ of every occupation $o \in \mathcal{O}$.[C.1]. Figure C.3 plots the slope of the corresponding regression line against simple measures that plausibly determine the size – the skill mean $\bar{s}$ for the task-shift effect and the skill dispersion S_s for the selection effect. Panel (a) shows that the task-shift effect increases faster in occupational exposure when $\bar{s}$ is smaller. The reason is that automating low-productivity tasks tends to lead to more powerful task-upgrading effects. The figure also indicates that the task-upgrading effect for tasks most exposed to AI ("Processing and Analyzing Records", "Reviewing and Editing Information") tend to be larger than those for skills exposed to robots ("Performing Detailed Manual Tasks", "Performing Physical Labor"). Panel (b) shows that selection effects typically put more downward pressure on wages when the associated skill is more dispersed. The reason is that tasks with more dispersion in skills across workers tend to be more dominant in determining sorting patterns and thereby exacerbate the correlation between occupational exposure and specialization in the automated task. Notably, the most affected tasks in the case of AI appear to be outliers from this pattern. Selection effects may therefore play a more important role in alternative scenarios of AI automation.

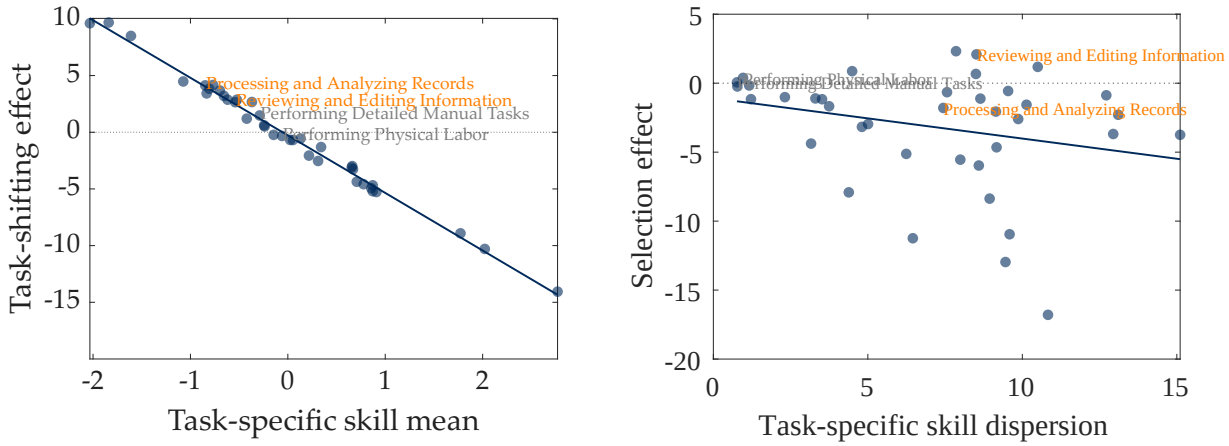
Alternative shock. Figure C.7 reproduces Figure 12 from the main text for an alternative shock to the second most affected task, "Maintaining and Managing Records" – also a form of information-processing. The results remain qualitatively unchanged.

C.2.3 Robustness: NLSY97

The baseline estimation is conducted using the National Longitudinal Survey of Youth 1997 (NLSY79). In this appendix, we show that the parameter estimates are not sensitive to using the National Longitudinal Survey of Youth 1997 (NLSY97) instead.

The NLSY97 is very similar to the NLSY79 in that it is one of the two National Longitudinal Surveys (NLS) administered by the Bureau of Labor Statistics. Many questions across the two surveys are identical and directly comparable. While the NLSY79 follows a cohort of about 6000 workers who entered the labor market in 1979, the NLSY97 follows a cohort of about 8000

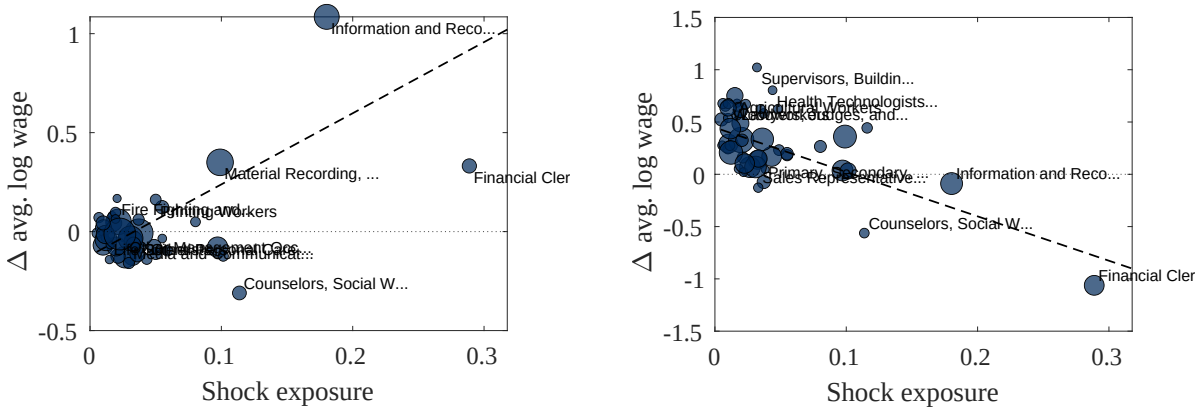
[C.1] We weight these regressions by pre-shock occupational employment shares



(a) Task-shifting **(b)** Selection

Figure C.3: Decomposition of task-shifting & selection effects: comparison across tasks

Notes. In both panels, each dot corresponds to a task. In the left panel, the horizontal axis displays the (standardized) estimated skill mean of the task, $\bar{s}_\tau$. The vertical axis displays the coefficient of a regression that regresses the magnitude of the task-shift effect from full automation for a given occupation on its exposure $A_{o,\tau}$, weighted by pre-shock employment. In the right panel, the horizontal axis displays the estimated between-worker dispersion of the skill associated with task τ: $S_{s,\tau}$. The vertical axis displays the coefficient of a regression that regresses the magnitude of the selection effect from full automation for a given occupation on its exposure $A_{o,\tau}$, weighted by pre-shock employment.



(a) Incumbent stayers **(b)** Incumbent leavers

Figure C.4: Stayers versus leavers

Notes. See notes for Figure 12.

workers who entered the labor market in 1997 and tracks their labor market outcomes over their lifetime. Given the later start date of the survey, the NLSY97 is shorter than the NLSY79 which is the reason why we prefer the NLSY79 as our baseline, given that identification comes in part from repeated observations associated with the same worker.

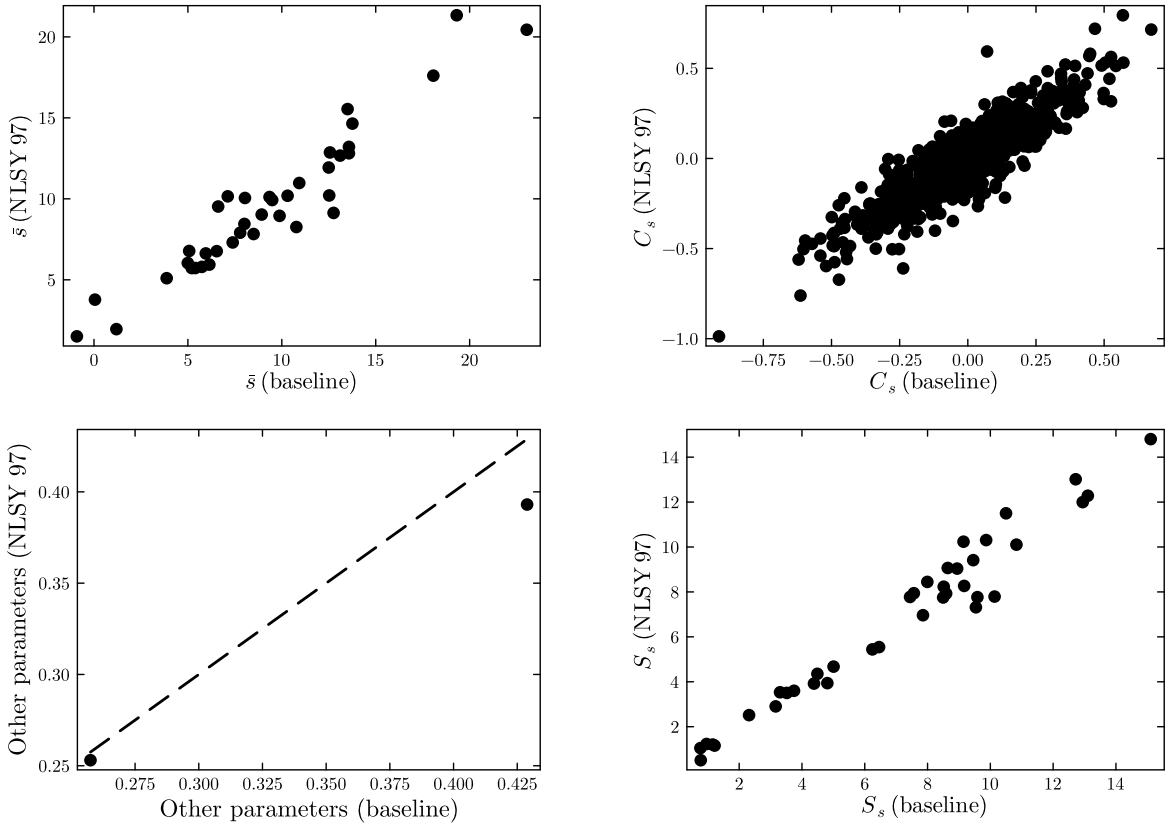


Figure C.5: Estimated parameters using the NLSY97 versus the baseline

Notes. The horizontal axis displays parameter values estimated in our baseline, which uses the NLSY79. The vertical axis displays corresponding estimated values when using the NLSY97. "Other parameters" refers to the tuple (ν, ς).

We repeat our estimation (holding A fixed) using the NLSY97. Figure C.5 shows the parameter estimates plotted against one another. The correlation of the parameter estimates across the two estimation versions is very high. We take this as evidence that (i) the skill distribution is robust across the two cohorts and has not changed substantially over time and (ii) our estimates are not driven by some features that are unique to our baseline data set.

C.2.4 Robustness: O*NET occupational task weights

In our baseline approach, we construct A based on the time shares inferred by an LLM, as explained in section 3.2. In this appendix, we explore an alternative estimate of A constructed from information about occupational task importance contained in the O*NET database. We

both directly compare the O*NET based A matrix to our baseline A matrix and discuss estimation and simulation results for this alternative A matrix.

Alternative A matrix. O*NET provides for each O*NET-SOC-2019 8-digit occupation not only a list of detailed task statements—from which we constructed our task clusters—but also assigns to each task a categorical rating on a Likert scale from 1-5 that indicates the importance of this task to its associated occupation.[C.2]

We begin aggregating up the importance ratings to the level of our aggregate task clusters by occupation. To do this, we first collapse occupations to the SOC-2019 minor group level. We then create weights for each occupation-cluster pair. These weights correspond to the shares of the detailed tasks associated to that occupation group that belong to a given cluster, weighted by importance ratings. Let k index occupations (at the minor-group level), τ index task clusters, and t index detailed tasks. Further, let $\mathcal{T}_\tau$ denote the set of detailed O*NET tasks associated with cluster τ, $\mathcal{T}_k$ the set of detailed tasks associated with occupation k, and $\omega_{k,t}$ be the weight attached to task τ by occupation k. Then we construct the (relative) weight occupation k puts on cluster τ, $\omega_{k,\tau}$ as

$$\omega_{k,\tau} = \frac{\sum_{t\in\mathcal{T}_\tau} \mathbf{1}_{t\in\mathcal{T}_k} \cdot \omega_{k,t}}{\sum_{t\in\mathcal{T}_k} \omega_{k,t}},$$

where $\mathbf{1}_{t\in\mathcal{T}_k}$ is an indicator function that equals 1 if task t belongs to occupation k and 0 otherwise. That is, the weight occupation k puts on cluster τ is greater if a large fraction of the detailed tasks associated to k are linked to τ or if those tasks have especially high importance weights for k. Next, the SOC-2019 occupations are cross-walked to the SOC-2000 classification used in our analysis using the official crosswalks available from https://www.onetcenter.org/taxonomy.html.

Comparing the occupational task weights thus obtained to our baseline, we find that they exhibit a strong, positive correlation (0.45). The O*NET task ratings yield an A matrix with a higher degree of concentration, in that the average occupation-level Herfindahl Index is 0.18 compared to 0.11 in our baseline.

Estimation and simulation results. To test to what extent our model results are robust to this alternative choice of A, we re-estimate the model with this alternative A. We then compare the implied skill distributions by contrasting the occupational wage of the median worker across both models. Figure C.6 shows that the mean occupational wages of the median worker across both sets of parameter estimates are highly correlated. In addition, the estimates for the standard deviations of idiosyncratic productivity and occupational preference shocks, ϱ (0.45) and ν (0.38), are not too far from the baseline (0.43 and 0.26), respectively.

[C.2] In addition to "importance," O*NET also provides scales for "relevance" and "frequency," which in principle could be used in the weighting also.

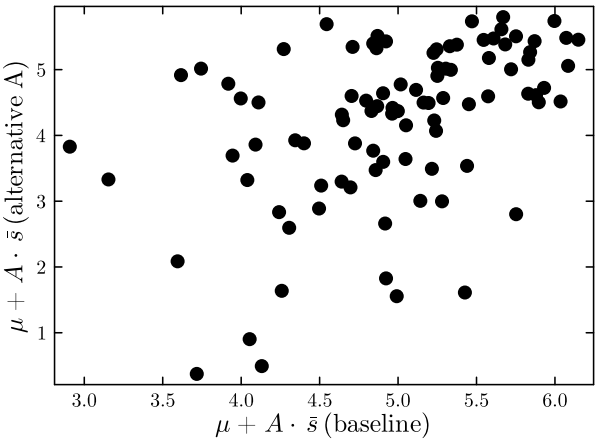


Figure C.6: Comparison of $\mu + A \cdot \bar{s}$ between baseline and alternative A

Notes. Each dot is an entry of $\mu + A \cdot \bar{s}$ given the baseline estimates (horizontal axis) or alternative estimates (vertical axis).

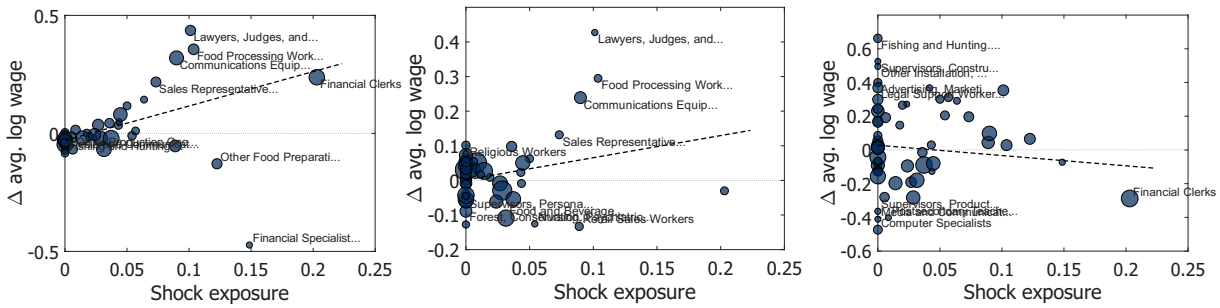
Lastly, we reproduce the thought experiment of fully automating "Processing and Analyzing Records" using the alternative parameter estimates. Specifically, Figure C.7 are the counterparts to the counterfactuals illustrated, in the baseline, in Figures 9 and 12. The results are qualitatively similar: Occupational wages rise more for more exposed occupations, more exposed incumbent stayers gain, and more exposed incumbent leavers lose. The magnitudes of the effects are notably smaller than in our baseline, though.

We conclude this section with a remark about why we prefer our baseline approach over the alternative way of constructing A outlined above. First, our baseline relies on cardinal time shares that identify the entries of the A matrix in a model-consistent way. Second, we find that the alternative models does a worse job in matching data moments such as average wages and occupational employment shares. Third, our baseline estimate is preferred through the lens of Bayesian model selection, as it is associated with a significantly higher likelihood.

C.2.5 Robustness: returns to occupational experience

As mentioned in the main text, a limitation of our baseline model is that it under-predicts workers' tendency to stay in the same occupation (cf. Figure 6). One likely reason is that the model abstracts from occupational tenure effects, which are important in the data (e.g., Traiberman, 2019, Figure 3). In this section, we we show that our model can be extended to feature a type of learning in the form of returns to occupational experience. Reassuringly, we found that all results are qualitatively — and indeed in most cases quantitatively — very similar to the baseline.

This version of the model is identical to our baseline except that a worker's productivity



(a) Occupational averages **(b)** Incumbent stayers **(c)** Incumbent leavers

Figure C.7: LLM automation – alternative estimates

Notes. Each dot is an occupation. The vertical axis of panel (a) measures the average occupational wage change before versus after the shock. The vertical axis of panel (b) measures the wage change for initial incumbents of each occupation before versus after the shock, including only those who stay in their occupation after the shock. The vertical axis of panel (c) measures the wage change for initial incumbents of each occupation before versus after the shock, including only those who leave their occupation after the shock. The horizontal axis of all panels measures shock exposure $A_{o,\tau}$. Dot sizes correspond to pre-shock employment shares. The dashed line is the line of best fit.

depends on their tenure in occupation o. Here, we consider the simplest such case, where productivity is greater whenever a worker has at least one year of experience in the occupation. We thus assume that workers who in period t choose occupation $o \in \mathcal{O}$ have a productivity that depends on whether they worked in occupation o in period $t-1$: If they did not work in o in $t-1$, their productivity is 1. If they did work in o in $t-1$, their productivity is $\exp(\Delta)$ with $\Delta > 0$

The worker's decision problem can be characterized as follows: Let

$$w^e_{i,o}(0) = \mu_o + A \cdot s_i$$
$$w^e_{i,o}(1) = \mu_o + \Delta + A \cdot s_i$$

be the expected wages of a worker with skills s_i. Then the worker's (expected) value function satisfies:

$$V_o(0) = w^e_{i,o}(0) + \beta\nu \log\left[\exp\left(\frac{V_o(1)}{\nu}\right) + \sum_{o' \neq o} \exp\left(\frac{V_{o'}(0)}{\nu}\right)\right]$$
$$V_o(1) = w^e_{i,o}(1) + \beta\nu \log\left[\exp\left(\frac{V_o(1)}{\nu}\right) + \sum_{o' \neq o} \exp\left(\frac{V_{o'}(0)}{\nu}\right)\right]$$

and thus that $V_o(1) = V_o(0) + \Delta$. This can be calculated via value function iteration and is simple enough to be handled by our estimation algorithm.

We do not report the full set of results for this specification but highlight that, as illustrated in

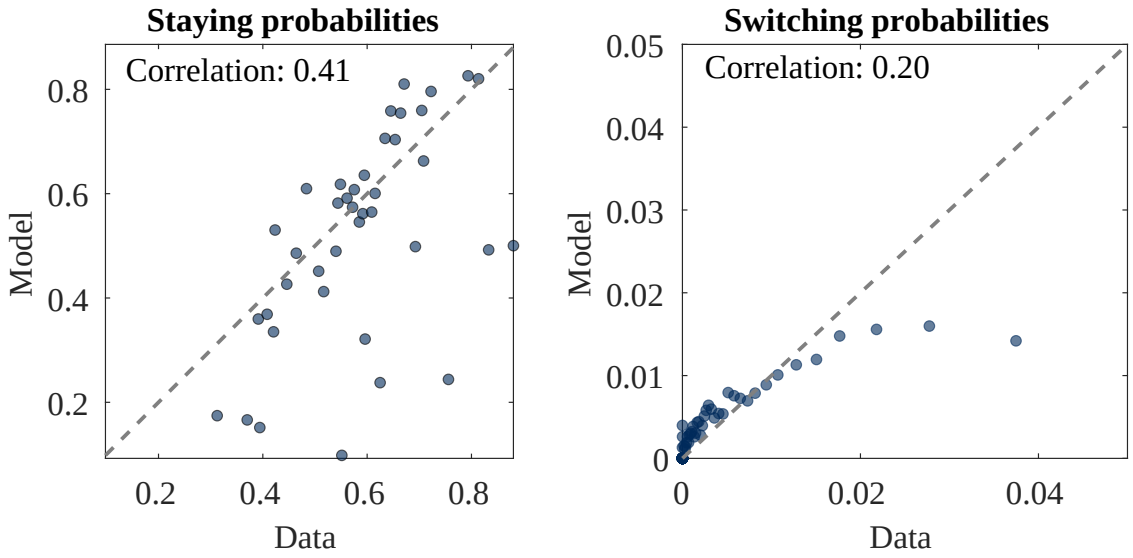


Figure C.8: Occupational transition patterns in the model with learning

Notes. See notes for Figure 6. For illustration, the model parameters are kept fixed at the baseline but we set $\Delta = 0.5$, yielding an average staying probability of 0.67, close to the empirical value of 0.63.

Figure C.8, this version of the model yields substantially greater levels of occupational persistence.