CES ifo

Working Papers

12779
2026

June 2026

The Price of Long Hours

Thimo De Schouwer, Sebastiaan Maes, Tom Potoms

CES ifo

Imprint:

CESifo Working Papers

ISSN 2364-1428 (digital)

Publisher and distributor: Munich Society for the Promotion
of Economic Research - CESifo GmbH

Poschingerstr. 5, 81679 Munich, Germany
Telephone +49 (0)89 2180-2740

Email office@cesifo.de
https://www.cesifo.org

Editor: Clemens Fuest

An electronic version of the paper may be downloaded free of charge
· from the CESifo website: www.ifo.de/en/cesifo/publications/cesifo-working-papers
· from the SSRN website: www.ssrn.com/index.cfm/en/cesifo/
· from the RePEc website: https://ideas.repec.org/s/ces/ceswps.html

The Price of Long Hours*

Thimo De Schouwer[1], Sebastiaan Maes[2], and Tom Potoms[3]

[1]KU Leuven
[2]University of Antwerp & CESifo
[3]University of Antwerp & IFS

June 28, 2026

Abstract

Earnings premiums for long hours vary across occupations, shaping gender sorting and inequality. The leading explanation emphasizes firm-side production technologies, yet hours and wages are equilibrium outcomes reflecting worker preferences and market structure. We build a directed search-and-matching model that allows us to decompose the premium. Premiums compensate for the disutility of long hours and bargaining power; productivity explains little in *Tech*, where premiums approach 20%. This disutility is highest among women, who sort away from long-hours jobs; these amenity differences account for 22% of the gender earnings gap. In counterfactuals, general-equilibrium responses through market tightness substantially change policy impacts.

Keywords: long hours, earnings premium, search and matching, directed search
JEL codes: C78, E24, J16, J31, J32

*This paper has benefited from comments and suggestions by Alex Bell, Christophe Bruneel-Zupanc, Liang Chen, Laurens Cherchye, Eugene Choo, Edoardo Ciscato, Pauline Corblet, Bart Cockx, Thibault Deneus, Jeremy Fox, Janoś Gabler, Marion Goussé, Iris Kesternich, Thomas Le Barbanchon, Tim Mensinger, Robert Miller, Amelie Schiprowski, Heiner Schumacher, Alessandra Voena, Jeff Wooldridge, and Hannah Zillessen, as well as audiences in Antwerp, Bergen (EALE), Birmingham (RES), Bonn (BGSE Applied Micro Workshop, IZA Summer School in Labor Economics), Bordeaux (EEA), Brussels (Belgian Day for Labour Economists), Leuven, Lisbon (IAAE), Luxemburg (Workshop on Gender and Economics), Palma (EWMES), Paris (COMPIE), and Seoul (ESWC). Paul Daniels and Marissa Hashizume provided generous assistance with the National Labor Exchange (NLx) Research Hub dataset. Thimo De Schouwer gratefully acknowledges financial support from the Research Foundation Flanders (FWO) through PhD Fellowship 11I7822N.

1 Introduction

There are substantial differences across occupations in the earnings premium associated with working long weekly hours.[1] As shown in Figure 1, U.S. premiums range from below 10% in some clinical occupations to more than 35% in some business occupations. These long-hours premiums are central to understanding gender inequality at the upper end of the earnings distribution (Blau & Kahn, 2017), since women are often less able to work long hours because they bear a larger share of home production responsibilities.[2] To explain why these premiums vary so sharply across occupations, Goldin (2014) and Goldin and Katz (2016) emphasize differences in how hours enter firms' production function. In some occupations, long hours substantially raise productivity, for example because client expectations or business models reward continuous availability, and this is reflected in large earnings premiums. In others, the productivity gains from longer hours are more limited, and the associated premium is correspondingly smaller.

Yet, as noted by Olivetti, Pan, and Petrongolo (2024), there is little empirical evidence establishing production technologies as the primary determinant of occupational long-hours premiums. Because hours and wages are equilibrium outcomes, they are shaped not only by firms' technologies, but also by workers' preferences and constraints, and by the market environment in which workers and firms interact. When working long weekly hours is more costly for workers, or when labor markets are tighter and workers enjoy stronger outside options, equilibrium long-hours premiums should adjust accordingly. Although we know little about the relative importance of these forces, disentangling technology from amenities is central to designing policies that reduce gender inequality, since the two call for very different responses.

This paper disentangles these forces by developing an empirically tractable search-and-matching model in which hours and wages are jointly determined. Our model embeds a directed-search environment as in Moen (1997), in which workers and firms direct their search ex ante toward partners with particular characteristics under specific hours arrangements, thereby generating local submarkets.[3] Within these submar-

[1] Bick, Blandin, and Rogerson (2022) document that a substantial share of total hours is supplied by individuals working more than 50 hours per week. A related literature finds sizable earnings premiums associated with such long hours; see, for example, Bertrand, Goldin, and Katz (2010); Goldin (2014); Cortés and Pan (2019); Mantovani (2023).

[2] See, for example, Aguiar and Hurst (2007); Adda, Dustmann, and Stevens (2017); Cortés and Pan (2019); Kleven, Landais, and Søgaard (2019).

[3] Directed search, in which workers and firms target specific submarkets rather than search randomly, provides a tractable and realistic description of the labor market. See Peters (1984); Moen (1997); Eeckhout and Kircher (2010, 2018) for seminal contributions and Wright, Kircher, Julien, and Guerrieri (2021) for a comprehensive overview.

Figure 1: Long-Hours Premiums by Occupation

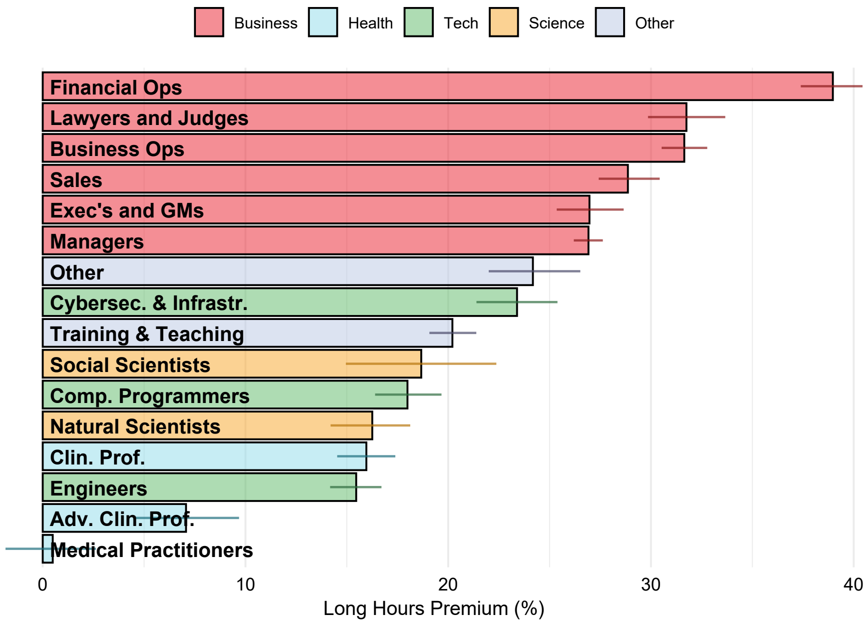

Notes. This figure shows the annual earnings premium associated with working long hours, by occupation, measured as the average percentage difference in earnings between long-hours jobs ($\geqslant$ 50 hours per week) and regular-hours jobs (35–49 hours per week). Estimates are based on data from the American Community Survey, pooled over the 2022–2024 waves. 95% confidence intervals are based on 200 bootstrap replications. The results are qualitatively similar when controlling for worker characteristics and region, and broadly consistent with those documented by Goldin (2014) and Cortés and Pan (2019).

kets, firms post wages that reflect both the productive value of a worker under a given hours arrangement and the probability of filling the vacancy, while workers choose among jobs on the basis of wages, amenities, and hiring probabilities. The interaction of these forces jointly determines equilibrium wages and hours. As a result, long-hours premiums arise endogenously, vary across workers and jobs, and admit a natural decomposition into three components: the utility costs of longer hours for workers, the productivity gains from longer hours for firms, and differences in market tightness across submarkets.[4]

We use this framework to address three questions. First, we assess the empirical relevance of the *Goldin hypothesis* by estimating occupation-specific productivity premiums for long hours. Second, we quantify the relative importance of productivity, disutility, and tightness across occupations. Third, we examine the implications for gender inequality, asking how each of these forces contributes to the gender earnings

[4] Long-hours premiums can be interpreted as compensating differentials (Rosen, 1986). In the presence of search frictions, however, they need not equal the marginal worker's valuation (Hwang, Mortensen, & Reed, 1998; Lang & Majumdar, 2004). In our model wages reflect a tightness-dependent split of surplus between workers and firms.

gap at the top of the distribution. We find that productivity premiums explain a smaller share of cross-occupation variation than the Goldin hypothesis would suggest, with worker disutility and market tightness playing substantial roles.

Our main methodological contribution is to introduce directed search into an equilibrium matching model with two-sided unobserved heterogeneity à la Choo and Siow (2006). We establish existence and uniqueness of a directed search equilibrium, in which realized market tightness coincides with agents' expectations, using results from Berry, Gandhi, and Haile (2013). Building on Chiappori, Salanié, and Weiss (2017) and Fox, Yang, and Hsu (2018), the model is identified from equilibrium matches and wages observed across multiple segmented labor markets. The key intuition is that, given the level of search frictions, matching patterns identify the joint surplus generated by a match. Observed wages then decompose this surplus into worker-side and firm-side components. Variation in the composition of workers and firms across labor markets identifies the search frictions themselves.

A distinguishing feature of our approach is that it only requires cross-sectional data at the aggregate submarket level. Labor economists have made substantial progress in understanding the role of non-wage amenities by exploiting job-to-job transitions in highly granular linked employer-employee data, through reduced-form (Card, Cardoso, & Kline, 2016; Sin, Stillman, & Fabling, 2022), structural (Sullivan & To, 2014; Taber & Vejlin, 2020; Lamadon, Lise, Meghir, & Robin, 2024; Xiao, 2024; Morchio & Moser, 2026), and combined approaches (Sorkin, 2018; Vattuone, 2026). Despite this growing attention to non-wage amenities, long hours have received comparatively little scrutiny, because detailed information on hours is generally available only through (repeated) cross-sectional surveys, which lack the richness needed for the linked worker-firm approaches above.

We estimate the model on U.S. survey data pooled over 2022–2024. Three main findings emerge. First, amenities matter more than productivity premiums in accounting for the long-hours premium, and bargaining power plays a substantial role in sectors such as *Tech* and *Health*. Second, our productivity estimates confirm the Goldin hypothesis for *Business* occupations but qualify it for *Tech*: there, sizable premiums stem not from productivity but from amenities and market tightness, as workers have a strong disutility of long hours while tight labor markets grant them considerable bargaining power. Third, the disutility of long hours is especially high among women, who therefore sort away from long-hours jobs and accept lower wages in regular-hours jobs. Gender differences in amenities account for 22% of the gender earnings gap. The model performs well out of sample. It closely reproduces the long-hours premium, which is not targeted in estimation, and likewise matches observed measures of market tightness. Reassuringly, our amenity estimates align with the survey-based estimates from the discrete-choice experiments of Maestas, Mullen, Powell, Von Wachter, and Wenger (2023), despite the two approaches drawing on entirely different data.

We conduct two counterfactual policy experiments to examine how the gender earnings gap responds to interventions on opposite sides of the labor market. The first operates on the demand side, eliminating the gender gap in how firms evaluate workers' productive value.[5] This perceived valuation gap accounts for 47% of the gender earnings gap directly, but closing it eliminates 69% once general-equilibrium forces—stronger relative demand for women, which tightens their labor markets and raises their bargaining power—are included. The second operates on the supply side, reallocating 20% of women from non-STEM to STEM degrees.[6] This reallocation narrows the gap by only about 1 percentage point, because the same tightness channel now runs in reverse: congestion in the submarkets women enter erodes their bargaining power and offsets much of the sorting gain. Taken together, the experiments deliver two lessons. First, general-equilibrium adjustment through market tightness can sharply alter, and even reverse, a policy's impact. Second, the gender gap stems mostly from how women's productive value is priced, not from how women sort across occupations.

Two strands of literature are most closely related to our paper. A first strand of related work studies how occupation-specific returns to long hours shape labor supply and gender gaps, taking those returns largely as primitives. Erosa, Fuster, Kambourov, and Rogerson (2022) provide a model-based quantitative assessment of the Goldin hypothesis, showing that gender differences in home-production responsibilities can account for a substantial share of gender gaps in occupational choice, wages, and hours.[7] Jang and Yum (2022) build a dynamic general-equilibrium model in which experience accumulates differently in nonlinear occupations, and use it to explain the rising share of women in such occupations and the narrowing of gender gaps in experience premiums over recent decades. We share these papers' emphasis on the importance of hours-related returns for sorting and gender inequality, but we ask a different question: rather than taking the cross-occupation pattern of long-hours premiums as given and studying the supply-side responses that follow, we treat those premiums as equi-

[5]These evaluations need not reflect actual productivity gaps by gender (e.g., from job tenure or job-specific human capital); they may equally capture statistical discrimination in how firms perceive men's and women's productivity, or taste-based discrimination.

[6]The gender gap in STEM is a contributing factor to the gender earnings gap, since STEM fields provide access to higher-paying occupations (Black, Haviland, Sanders, & Taylor, 2008; Blau & Kahn, 2017; Sloane, Hurst, & Black, 2021). A growing literature evaluates policy interventions aimed at increasing women's entry into STEM (Carrell, Page, & West, 2010; Breda, Grenet, Monnet, & van Effenterre, 2023).

[7]A related analysis is Erosa, Fuster, Kambourov, and Rogerson (2024), who study wage and earnings inequality between and within occupations and show that the interplay between heterogeneous tastes for leisure and occupational nonlinearities in the hours-earnings relationship is central to explaining the joint distribution of hours, wages, and earnings.

librium objects and examine where they come from. This requires moving beyond competitive sorting to a frictional environment in which firms post wages and workers direct their search, so that productivity, disutility, and market tightness can be separately identified as components of the premium.

A second strand uses equilibrium search models to recover non-wage amenities and decompose the gender wage gap. Flabbi and Moro (2012) develop a search-and-bargaining model with wages and work-hours flexibility, and show that a sizable share of women place positive value on flexibility such that reducing its cost would narrow the gender wage gap. Closest to us in objective, Morchio and Moser (2026) estimate an equilibrium search model with endogenous firm pay, amenities, and hiring on linked employer-employee data, and find that compensating differentials explain roughly half of the gender pay gap. Their focus is on firm-level amenities broadly construed rather than on (long) hours specifically. We share with this strand the goal of recovering non-wage attributes from equilibrium outcomes, but differ in market structure: by embedding directed search rather than random search with bargaining, vacancy posting responds endogenously and submarket tightness becomes a third determinant of equilibrium wages alongside productivity and amenities. This margin is silent in Flabbi and Moro (2012). It maps onto, but is conceptually distinct from, the within-firm bargaining channel in Morchio and Moser (2026). In our model the surplus split is a property of the submarket rather than of the worker-firm pair, with tightness operating as a market-level margin that the within-firm focus of that paper leaves implicit.[8]

Outline of the paper. The remainder of this paper is structured as follows. Section 2 introduces our equilibrium directed search model in which wages and hours are jointly determined. Section 3 discusses identification and structural estimation, using equilibrium matches and wages observed across multiple segmented labor markets. Section 4 describes the data. Section 5 presents our productivity and amenity estimates and the associated earnings decompositions. Section 6 reports the findings of two counterfactual policy experiments. Section 7 concludes. All proofs and additional results are in the Appendix.

2 Directed Search with Endogenous Job Attributes

To understand what shapes long-hours premiums, we develop a directed search and matching framework with two-sided heterogeneity in which wages and hours are jointly endogenous.

[8] By taking a market-level perspective, we abstract from micro-level sources of bargaining such as gender differences in salary negotiation (Bertrand, 2011; Leibbrandt & List, 2015).

2.1 Setup

Consider a labor market populated by workers $i \in \mathcal{I}$ and firms $j \in \mathcal{J}$. Workers and firms are grouped into mutually exclusive discrete types $x \in \mathcal{X}$ and $y \in \mathcal{Y}$, respectively. Let $N_x > 0$ and $N_y > 0$ denote the number of workers and firms of types x and y. Jobs are characterized by a set of mutually exclusive discrete attributes $z \in \mathcal{Z}$, capturing non-pecuniary aspects of employment. In our setting, it is useful to interpret $\mathcal{Z}$ as the set of possible hours arrangements. More generally, however, these attributes may reflect features such as workplace flexibility or in-kind benefits.

The search process is not random, but directed. In particular, we assume that both workers and firms meet each other in endogenously chosen submarkets, indexed by a triple (x, y, z). In such a submarket, workers of type x are searching for firms of type y who offer hours z. Likewise, firms of type y are searching for workers of type x who seek hours z. In each submarket, workers match to firms one-to-one or remain unassigned.[9] Before further expanding on this, we will first describe the expected utility and profit of workers and firms, respectively.

Expected utility. The expected utility of a worker i of type x receiving a wage W in submarket (x, y, z) is given by

$$U_{x_iyz} \equiv P(\theta_{xyz})(\alpha_{xyz} + W) + \varepsilon_{x_iyz}. \tag{1}$$

Here, α_{xyz} captures the non-monetary amenities that are derived from a match. These reflect a combination of preferences, social norms, and non-market constraints, which our framework cannot separately identify.[10] $P(\theta_{xyz})$ denotes the job-finding probability, which is a function of the degree of (sub-)market tightness $\theta_{xyz} > 0$. Tightness is defined as the ratio of the number of vacancies and the number of searchers in the market. Consequently, the first term gives the expected utility from searching, i.e., the probability of finding a job multiplied by the total utility of a job in this submarket.[11] Finally, the random disturbance ε_{x_iyz} reflects idiosyncratic differences in search costs between workers. Such a disturbance rationalizes that observationally identical workers sort into different types of jobs and hours arrangements.

Following Menzio and Shi (2010), we specify the job finding probability as a constant elasticity of substitution (CES) function of market tightness:

$$P(\theta_{xyz}) \equiv \left(1 + \theta_{xyz}^{\lambda}\right)^{\frac{1}{\lambda}}. \tag{2}$$

[9] In one-to-one matching models, firms and jobs are interchangeable. This is equivalent to firms hiring multiple workers if the output is separable across jobs, see for example Roth and Sotomayor (1989).

[10] Treating the amenity term as such a composite is common in the literature; see, for example, Le Barbanchon, Rathelot, and Roulet (2021); Morchio and Moser (2026).

[11] We normalize the value of remaining unmatched to zero for both workers and firms.

Parameter $\lambda < 0$ is the elasticity of substitution and captures the curvature of the job finding function, and thereby shapes the complementarity of the matching process.

Expected profit. Similarly, the expected profit of a firm j of type y paying a wage W in submarket (x, y, z) is given by

$$\Pi_{xy_jz} \equiv Q(\theta_{xyz})(\rho_{xyz} - W) + \eta_{xy_jz}. \tag{3}$$

Here, ρ_{xyz} captures the productive value of a match. It reflects a combination of true match-specific productivity, firms' perceptions or statistical inferences about productivity, and any taste-based component in how firms value an x-type worker with hours arrangement given at z. The term $Q(\theta_{xyz}) \equiv \frac{P(\theta_{xyz})}{\theta_{xyz}}$ denotes the hiring probability. Consequently, the first term gives the expected profit from searching, i.e., the probability of finding a worker multiplied by the productive value net of wages. Finally, the random disturbance η_{xy_jz} captures idiosyncratic variation in screening costs across firms.

Equation (3) captures a key feature of our model. Even when two worker types x and x' are equally productive ($\rho_{xyz} = \rho_{x'yz}$), firms may earn different expected profits from matching with them, because the hiring probability depends on the tightness of the submarket. As we show below, this generates equilibrium wage differences across submarkets even at equal productivity: firms can pay lower wages in slacker markets, where workers face worse outside options.

2.2 Directed Search Equilibrium

We now turn to a description of how workers and firms sort themselves in submarkets, and formally discuss our equilibrium notion.

Wages. Wages are the central force allocating agents across submarkets: holding all else fixed, a higher posted wage draws workers in and deters firms, shifting market tightness. A firm's wage and the tightness it faces are therefore jointly determined, and in principle the firm would optimize over both. Following the competitive-search literature (Montgomery, 1991; Moen, 1997; Wright et al., 2021), we instead assume each firm is atomistic and posts wages taking as given the market utility that workers can obtain by searching elsewhere. A single firm's posting therefore does not move this market utility. Formally, firms in submarket (x, y, z) solve

$$\begin{gathered} \max_{W} \; Q(\theta_{xyz})(\rho_{xyz} - W), \\ \text{subject to } P(\theta_{xyz})(\alpha_{xyz} + W) = U_x, \end{gathered} \tag{4}$$

where U_x is the market utility of a worker of type x.[12] The participation constraint in Equation (4) then implicitly defines a map $W_{xyz}(\theta_{xyz})$ between wages and tightness in the submarket. Substituting this relation into the objective and solving the first-order condition with respect to θ_{xyz} yields the following expression for wages in this submarket:[13]

$$W_{xyz}(\theta_{xyz}) = (1 - F(\theta_{xyz}))\rho_{xyz} - F(\theta_{xyz})\alpha_{xyz}. \tag{5}$$

Here, $F(\theta_{xyz}) \equiv \theta_{xyz}\frac{P'(\theta_{xyz})}{P(\theta_{xyz})} = \frac{\theta_{xyz}^{\lambda}}{1+\theta_{xyz}^{\lambda}}$, which is the elasticity of the job-finding probability with respect to market tightness. The term $F(\theta_{xyz})$ can be interpreted as the firm's bargaining power, while $1 - F(\theta_{xyz})$ corresponds to the worker's bargaining power.[14] Indeed, the greater the worker's bargaining power, the larger the share of productive value ρ_{xyz} that accrues to the worker.

Equilibrium wages in Equation (5) increase in productivity ρ_{xyz} due to competition between firms, and decreases in the cost of work α_{xyz} due to competition between workers. But the extent to which these forces matter is governed by tightness in the labor market, through the matching function elasticity $F(\theta_{xyz}) \in (0, 1)$. The elasticity is larger when there is slack in the market ($\theta_{xyz} \to 0$), and firms can exploit their superior bargaining positions to capture a larger share of the joint surplus.

To better understand the role of market tightness, note that in the slack limit when $\theta_{xyz} \to 0$, we have $F(\theta_{xyz}) \to 1$ at interior values of the substitution elasticity. This means that firms post wages to only just reward workers for their cost of employment, at exactly $-\alpha_{xyz}$. When markets tighten, $F(\theta_{xyz})$ decreases, as many firms compete for a limited number of workers. This pushes up wages, and in the limit when $\theta_{xyz} \to \infty$ and $F(\theta_{xyz}) \to 0$, it allows workers to capture the entire value of production ρ_{xyz}.[15]

The elasticity of substitution λ governs how responsive the perceived bargaining weights $F(\theta_{xyz})$ are to changes in market tightness. In the Leontief limit when $\lambda \to -\infty$, we have $F(\theta_{xyz}) \in \{0, 1\}$, and wages are set such that the short side of the market captures the entire match surplus. In the Cobb-Douglas limit when $\lambda \to 0$, we have $F(\theta_{xyz}) = 0.5$, and wages always split the surplus evenly, regardless of the market conditions. At interior values, the perceived bargaining weights are a smooth function

[12] The idiosyncratic search and screening costs do not play a role in the determination of wages as they are sunk costs once workers and firms have entered a submarket.

[13] In this expression, market utility does not appear since its role is fully absorbed by tightness: a higher U_x requires a higher θ_{xyz} to satisfy the participation constraint, and conditional on θ_{xyz}, the wage is fully determined.

[14] As pointed out by Wright et al. (2021), this approach to obtain equilibrium wages is equivalent to workers and firms Nash bargaining over the match surplus, where the bargaining weight is given by the job finding elasticity. This is essentially equivalent to the well-known Hosios (1990) efficiency condition.

[15] This assumes throughout that $\alpha_{xyz} + \rho_{xyz} > 0$ such that there are systematic gains from matching, and maintains that $\lambda < 0$.

of market tightness.

Before turning to directed search, note that substituting the wage Equation (5) into Equations (1) and (3) allows expected utility and profits to be written as functions of market tightness:

$$\begin{aligned} U_{x_i yz} &= \underbrace{P(\theta_{xyz})(1 - F(\theta_{xyz}))\Phi_{xyz}}_{\equiv \overline{U}_{xyz}(\theta_{xyz})} + \varepsilon_{x_i yz}, \\ \Pi_{xy_j z} &= \underbrace{Q(\theta_{xyz})F(\theta_{xyz})\Phi_{xyz}}_{\equiv \overline{\Pi}_{xyz}(\theta_{xyz})} + \eta_{xy_j z}, \end{aligned} \tag{6}$$

where $\Phi_{xyz} \equiv \alpha_{xyz} + \rho_{xyz} > 0$ denotes the total surplus generated by an (x, y, z)-match. Lemma A.1 in the Appendix implies that, ceteris paribus,

$$\frac{\partial U_{x_i yz}}{\partial \theta_{xyz}} > 0, \quad \text{and} \quad \frac{\partial \Pi_{xy_j z}}{\partial \theta_{xyz}} < 0,$$

This is intuitive: greater market tightness benefits the bargaining position of workers (and should therefore increase their expected utility) while harming the position of firms (and should therefore reduce their expected profit).

Unobserved heterogeneity. The disturbances $\varepsilon_{x_i yz}$ and $\eta_{xy_j z}$, reflecting idiosyncratic search and screening costs, are assumed to be independently and identically distributed according to a Type I Extreme Value distribution. This allows us to express optimal choice behavior of workers and firms types in closed-form (McFadden, 1978).

To keep the analysis tractable, we impose separability of unobservables on both sides of the market.[16] Specifically, we assume that worker-side shocks $\varepsilon_{x_i yz}$ are independent of which particular firm of a given type a worker matches with, and symmetrically that firm-side shocks $\eta_{xy_j z}$ are independent of the particular workers within a given type. This is a strong assumption, as it rules out that interactions between the unobservables on both sides of the market create additional surplus. In our empirical application, we mitigate this concern in two ways. First, we restrict the sample to a relatively homogeneous group by selecting only highly productive individuals from the set of occupations defined in Goldin (2014). Second, we include rich observable type classifications that capture a substantial share of systematic heterogeneity in the joint match surplus.

Choice over submarkets. Optimizing expected utility and profit in Equation (6) leads to a simultaneous two-sided discrete choice model. Conditional on their beliefs regarding market tightness, both sides of the market search in a directed way over potential

[16]This is a common assumption in econometric models of two-sided matching (Choo & Siow, 2006; Dupuy & Galichon, 2014). For applications on matching in the labor market, see, for example, Dupuy and Galichon (2022) and Caldwell and Danieli (2024).

partners and hours. Workers search on the submarket that maximizes their expected utility. The proportion of workers of type x searching for firms of type y with hours z is given by

$$\pi^w_{yz|x} \equiv \Pr_{\varepsilon}\left[U_{x_i yz}(\theta_{xyz}) \geqslant \{U_{x_i y'z'}(\theta_{xy'z'})\}_{y',z'}\right] = \frac{\exp\left(\overline{U}_{xyz}(\theta_{xyz})\right)}{\sum_{y',z'}\exp\left(\overline{U}_{xy'z'}(\theta_{xy'z'})\right)}. \tag{7}$$

Likewise, firms search on the submarket that maximizes their expected profits. The proportion of firms of type y searching for workers of type x with hours z is given by

$$\pi^f_{xz|y} \equiv \Pr_{\eta}\left[\Pi_{xy_j z}(\theta_{xyz}) \geqslant \{\Pi_{x'y_j z'}(\theta_{x'yz'})\}_{x',z'}\right] = \frac{\exp\left(\overline{\Pi}_{xyz}(\theta_{xyz})\right)}{\sum_{x',z'}\exp\left(\overline{\Pi}_{x'yz'}(\theta_{x'yz'})\right)}. \tag{8}$$

From Equations (7) and (8), we can derive the total number of searchers on each side of submarket (x, y, z):

$$S^w_{yz|x}(\boldsymbol{\theta}) \equiv \pi^w_{yz|x} N_x, \quad \text{and} \quad S^f_{xz|y}(\boldsymbol{\theta}) \equiv \pi^f_{xz|y} N_y.$$

Ignoring general-equilibrium effects, Lemma A.2 in the Appendix implies that

$$\frac{\partial S^w_{yz|x}}{\partial \theta_{xyz}} > 0, \quad \frac{\partial S^w_{yz|x}}{\partial \theta_{xy'z'}} < 0, \quad \frac{\partial S^f_{xz|y}}{\partial \theta_{xyz}} < 0, \quad \text{and} \quad \frac{\partial S^f_{xz|y}}{\partial \theta_{x'yz'}} > 0,$$

where $(x', y', z') \neq (x, y, z)$. That is, the mass of workers of type x who choose to search in submarket (x, y, z) is increasing in the tightness of that submarket. For firms, the reverse holds: the mass of firms of type y that choose to search in the same submarket is decreasing in its tightness. Moreover, an increase in tightness in another submarket (x, y', z') reduces workers' incentives to search in (x, y, z), with the reverse conclusion again applying to firms.

Equilibrium tightness. For the beliefs of workers and firms over market tightness θ_{xyz} to be self-supporting, in equilibrium those beliefs must coincide with realized tightness. This leads to our formal notion of *directed search equilibrium*. An attractive property of this equilibrium notion is that it does not require markets to clear, hence allowing for frictional unemployment.

Definition 1. *A directed search equilibrium is a market tightness vector $\boldsymbol{\theta}^* \in \Theta \subseteq \mathbb{R}_+^{|\mathcal{X}|\times|\mathcal{Y}|\times|\mathcal{Z}|}$ that solves the $|\mathcal{X}| \times |\mathcal{Y}| \times |\mathcal{Z}|$-dimensional system of equations:*

$$\theta_{xyz} = \frac{S^f_{xz|y}(\boldsymbol{\theta})}{S^w_{yz|x}(\boldsymbol{\theta})}, \quad \forall x, y, z. \tag{9}$$

The equilibrium is unique if there exists only one such $\boldsymbol{\theta}^*$.

The directed search equilibrium exists under the maintained assumptions and is unique under an additional condition, typically satisfied when worker and firm types are sufficiently heterogeneous. Together, these results ensure *coherency* and *completeness* of the structural model in the sense of Tamer (2003).[17]

Proposition 1. *A directed search equilibrium exists.*

Existence follows from applying Brouwer's fixed point theorem to Equation (9): the equilibrium map is a continuous self-map on a compact convex domain, so a fixed point always exists.

Proposition 2. *Suppose there exists a submarket* (x, y, z) *such that, for all* $\boldsymbol{\theta} \in \Theta$,

$$\pi^{w}_{yz|x}(\boldsymbol{\theta}) \leqslant \frac{1}{|\mathcal{Y}||\mathcal{Z}|}, \quad \text{and} \quad \pi^{f}_{xz|y}(\boldsymbol{\theta}) \leqslant \frac{1}{|\mathcal{X}||\mathcal{Z}|}. \tag{10}$$

Then the directed search equilibrium is unique.

The proof verifies the connected weak substitutes property of Berry et al. (2013). It suffices that there exists a submarket that is unattractive to *both* sides of the market relative to the other submarkets a worker or firm of the same type could choose. The condition is mild in two respects. First, it only needs to hold for a *single* submarket. Second, the thresholds have a natural interpretation as the uniform-search benchmark. For example, $1/(|\mathcal{Y}||\mathcal{Z}|)$ is the probability with which a worker of type x would direct search to (y, z) if indifferent across all options. The condition therefore only requires that the submarket be no more attractive than indifference.

By Jensen's inequality, a sufficient condition for the inequalities in Equation (10) to hold is that expected utility and expected profit each fall below their within-type averages in some submarket (x, y, z):

$$\overline{U}_{xyz}(\theta_{xyz}) \leqslant \frac{1}{|\mathcal{Y}||\mathcal{Z}|}\sum_{y',z'} \overline{U}_{xy'z'}(\theta_{xy'z'}), \quad \text{and} \quad \overline{\Pi}_{xyz}(\theta_{xyz}) \leqslant \frac{1}{|\mathcal{X}||\mathcal{Z}|}\sum_{x',z'} \overline{\Pi}_{x'yz'}(\theta_{x'yz'}),$$

for all $\boldsymbol{\theta} \in \Theta$.[18] Such "bad corners" are typically present under sufficient heterogeneity in worker and firm types: generically some submarket is less attractive to both sides than the alternatives.

Equilibrium matches and wages. Market tightness is typically unobserved by the researcher, since it requires information on the submarkets in which unmatched workers search. It is therefore useful to characterize how the directed search equilibrium maps into observable outcomes, namely equilibrium matches M_{xyz} and wages W_{xyz}.

[17] For every admissible set of primitives, the model delivers exactly one equilibrium outcome: existence guarantees coherency, uniqueness guarantees completeness.

[18] This is shown in Lemma A.5 in the Appendix.

These objects will play a central role in Section 3, where we discuss identification and estimation of the model parameters.

Equilibrium matches follow directly from the directed search equilibrium. In particular, the number of realized matches in submarket (x, y, z) is given by

$$M_{xyz} \equiv P(\theta^*_{xyz}) S^w_{yz|x}(\boldsymbol{\theta}^*) = Q(\theta^*_{xyz}) S^f_{xz|y}(\boldsymbol{\theta}^*). \tag{11}$$

This expression reflects that the number of matches equals the number of agents (workers or firms) entering the submarket multiplied by the probability of a successful match. The two equalities follow from market clearing in each submarket. Similarly, the number of unmatched agents is given by

$$\begin{aligned} M_{x0} &\equiv \sum_{y,z} (1 - P(\theta^*_{xyz})) S^w_{yz|x}(\boldsymbol{\theta}^*), \\ M_{0y} &\equiv \sum_{x,z} (1 - Q(\theta^*_{xyz})) S^f_{xz|y}(\boldsymbol{\theta}^*). \end{aligned} \tag{12}$$

Here, M_{x0} denotes the number of unemployed workers of type x, and M_{0y} the number of firms of type y (vacancies) that remain unmatched.[19] From Equation (5), we have that equilibrium wages are given by

$$W_{xyz} \equiv (1 - F(\theta^*_{xyz}))\rho_{xyz} - F(\theta^*_{xyz})\alpha_{xyz}. \tag{13}$$

Across all submarkets, we stack equilibrium matches and wages in vectors $\mathbf{M}$ and $\mathbf{W}$, respectively.

3 Empirical Implementation

We now turn to identification and estimation. Following Chiappori et al. (2017) and Fox et al. (2018), we assume that the researcher observes equilibrium matches and wages across multiple segmented markets $s \in \mathcal{S}$.[20] In each market, the researcher can construct consistent estimates $\widehat{\mathbf{M}}^s$ and $\widehat{\mathbf{W}}^s$. For notational convenience, we stack these objects into a single data vector, $\widehat{\mathbf{D}} \equiv \left(\widehat{\mathbf{M}}^1, \widehat{\mathbf{W}}^1, \ldots, \widehat{\mathbf{M}}^{|\mathcal{S}|}, \widehat{\mathbf{W}}^{|\mathcal{S}|}\right)$. We denote the associated population object as $\mathbf{D}$.

[19]Note that number of unmatched workers and firms, given by Equation (12), do not give additional independent information for identification as the accounting identities

$$M_{x0}(\boldsymbol{\theta}) + \sum_{y,z} M_{xyz}(\boldsymbol{\theta}) = N_x, \qquad M_{0y}(\boldsymbol{\theta}) + \sum_{x,z} M_{xyz}(\boldsymbol{\theta}) = N_y,$$

must always hold.

[20]Note that the model is not identified from data on a single labor market, since there are $2|\mathcal{X}||\mathcal{Y}||\mathcal{Z}|+1$ unknowns but only $2|\mathcal{X}||\mathcal{Y}||\mathcal{Z}|$ equations. As we discuss below, however, observing even one additional market already generates substantial overidentification.

Importantly, the researcher does not observe equilibrium market tightness, since the search decisions of unmatched agents are unobserved. This constitutes a central challenge, as market tightness is a key determinant of both match surplus and job-finding probabilities. Despite this limitation, we demonstrate that variation in the marginal distribution of types across markets suffices to identify the model primitives and, through the structure of the model, equilibrium market tightness. We assume that all markets share the same surplus structure.

3.1 Discussion on Identification

The model is parameterized by submarket-specific amenity parameters (α_{xyz}) and productivity parameters (ρ_{xyz}), and the matching-function parameter (λ). Let $\boldsymbol{\beta}$ collect these model primitives. We assume that $\boldsymbol{\beta}$ is common across markets and exploit variation in the composition of worker types (N_x) and firm types (N_y) across markets for identification. Such compositional differences shift market tightness and, in turn, induce variation in equilibrium matching patterns and wages. This variation is used to recover $\boldsymbol{\beta}$.

The key intuition for identification proceeds in three steps. Suppose for the moment that search frictions, captured by λ, are known. First, the observed number of matches in submarket (x, y, z) then reveals information about the total match surplus Φ_{xyz}: submarkets with higher surplus attract more searchers on both sides, and therefore generate more matches. Second, given this surplus, the observed wage W_{xyz} separates Φ_{xyz} into its worker and firm components, α_{xyz} and ρ_{xyz}. Holding total surplus fixed, a higher wage implies that the amenities α_{xyz} must be lower, and conversely. This is the standard compensating-differentials logic. Third, variation across multiple submarkets provides the information needed to identify λ itself. In the Leontief limit when $\lambda \to -\infty$, small surplus differences generate sharp differences in matching volumes across submarkets, as most agents pile into those with high surplus. In the Cobb-Douglas limit when $\lambda \to 0$, the same surplus differences produce more muted variation in matches across submarkets.

Because the model is nonlinear and incorporates general-equilibrium interactions, a full identification result is difficult to characterize analytically. The example below nonetheless illustrates that $\boldsymbol{\beta}$ is generically identified, provided that labor market composition varies sufficiently across markets.

Example 1. *Consider two labor markets, denoted by s and s', each populated by a single worker type x and a single firm type y. The corresponding population counts are $N_x^s = 10$, $N_y^s = 15$, $N_x^{s'} = 10$, and $N_y^{s'} = 20$, so the two markets contain the same number of workers but s' has a third more firms. Suppose that there are two possible work arrangements, regular hours z and long hours z', and let $\boldsymbol{\beta}_0 \equiv (\alpha_{xyz}, \alpha_{xyz'}, \rho_{xyz}, \rho_{xyz'}, \lambda) = (2, 0, 0, 1, -1)$ denote the true* *parameter vector. These parameter values imply that regular hours are associated with higher amenities, but lower productive value, than long hours.*

Calculations show that the Jacobian of $\mathbf{D}$ *with respect to* $\boldsymbol{\beta}$*, evaluated at* $\boldsymbol{\beta}_0$*, is given by*

$$\frac{\partial \mathbf{D}(\boldsymbol{\beta}_0)}{\partial \boldsymbol{\beta}} = \begin{array}{c} \\ \partial M^s_{xyz} \\ \partial M^s_{xyz'} \\ \partial W^s_{xyz} \\ \partial W^s_{xyz'} \\ \partial M^{s'}_{xyz} \\ \partial M^{s'}_{xyz'} \\ \partial W^{s'}_{xyz} \\ \partial W^{s'}_{xyz'} \end{array} \begin{array}{c} \begin{array}{ccccc} \partial \alpha_{xyz} & \partial \alpha_{xyz'} & \partial \rho_{xyz} & \partial \rho_{xyz'} & \partial \lambda \end{array} \\ \left[\begin{array}{rrrr|r} 0.381 & -0.389 & 0.381 & -0.389 & -2.558 \\ -0.389 & 0.401 & -0.389 & 0.401 & -1.468 \\ -0.431 & 0.026 & 0.569 & 0.026 & -0.156 \\ 0.016 & -0.404 & 0.016 & 0.596 & -0.124 \\ \hline 0.449 & -0.475 & 0.449 & -0.475 & -2.775 \\ -0.475 & 0.510 & -0.475 & 0.510 & -1.432 \\ -0.384 & 0.041 & 0.616 & 0.041 & -0.249 \\ 0.025 & -0.340 & 0.025 & 0.660 & -0.193 \end{array}\right] \end{array}.$$

The upper and lower blocks correspond to labor markets s *and* s'*, respectively. Within each market, the first two rows give the derivatives of matches, and the last two rows the derivatives of wages. Since the Jacobian matrix has full column rank,* $\boldsymbol{\beta}$ *is locally identified in an open neighborhood of* $\boldsymbol{\beta}_0$*.*[21] *The full-rank property reflects the three-step intuition developed above:*

1. *Matches identify the joint surplus, but not its components.*
   In the match rows, the columns for α_{xyz} *and* ρ_{xyz} *are identical, and likewise for* $\alpha_{xyz'}$ *and* $\rho_{xyz'}$*. This reflects that matches depend on amenities and productivity only through the total match surplus. Looking at the* 2×4 *matches block in isolation, the* α *and* ρ *columns are perfectly collinear: observed matches alone may not disentangle both components. The joint surplus itself, however, is identified conditional on the unknown parameter* λ*.*

2. *Wages allow to decompose the joint surplus.*
   In the wage row W_{xyz}*, the columns for* α_{xyz} *and* ρ_{xyz} *have opposite signs, and likewise for row* $W_{xyz'}$ *and the columns for* $\alpha_{xyz'}$ *and* $\rho_{xyz'}$*. This reflects the compensating-differentials logic. The entire* 4×4 *left-hand block therefore has full column rank: the wage rows break the* α*-*ρ*-collinearity in the matches block, so that* α_{xyz}*,* $\alpha_{xyz'}$*,* ρ_{xyz}*, and* $\rho_{xyz'}$ *are jointly identified within each market, conditional on the unknown parameter* λ*.*

3. *Variation in market composition identifies the matching function elasticity.*
   The difference in the number of firms across s *and* s' *shifts equilibrium tightness, and* λ *governs how strongly those shifts translate into matching and wage outcomes. Non-proportional rows across* s *and* s' *are the observable consequence, and provide the information that pins down* λ*.*

[21] If $\frac{\partial \mathbf{D}(\boldsymbol{\beta}_0)}{\partial \boldsymbol{\beta}}$ has full column rank, any nonzero local perturbation of $\boldsymbol{\beta}$ around $\boldsymbol{\beta}_0$ induces a change in $\mathbf{D}(\boldsymbol{\beta})$. Our argument is related to the notion of local identification (Rothenberg, 1971; Lewbel, 2019).

3.2 Estimation

To estimate the model, we combine a minimum-distance estimator with a nested root-finding algorithm. The outer loop optimizes over the parameter space, and for each candidate parameter vector the inner loop solves for the directed-search equilibrium required to compute equilibrium matches and wages.

Minimum-distance estimator. Based on the observable implications of the model, we estimate $\boldsymbol{\beta}$ by minimizing a loss function that penalizes the overall (weighted) distance between $\mathbf{D}(\boldsymbol{\beta})$ and $\widehat{\mathbf{D}}$. In the outer loop, we solve

$$\widehat{\boldsymbol{\beta}} \equiv \arg\min \left(\widehat{\mathbf{D}} - \mathbf{D}(\boldsymbol{\beta})\right)^{\intercal} \boldsymbol{\Omega} \left(\widehat{\mathbf{D}} - \mathbf{D}(\boldsymbol{\beta})\right),$$

where $\boldsymbol{\Omega}$ is a positive definite weighting matrix. This estimator is consistent and asymptotically normal under standard regularity conditions.[22] The optimal weighting matrix that produces the smallest variance is equal to the inverse of the variance-covariance matrix of the moments. We obtain an empirical estimate of this variance-covariance matrix through a bootstrapping procedure. Standard errors for the parameters are obtained analytically.

In the inner loop, we compute $\mathbf{D}(\boldsymbol{\beta})$ for each candidate parameter vector. We begin by solving Equation (9) using the bisection method, a robust root-finding procedure. We then use the resulting equilibrium market tightness to compute matching patterns and wages via Equations (11) and (13).

Parameterization. To improve interpretability and mitigate the curse of dimensionality, we adopt a flexible yet parsimonious specification for both utility and profit parameters. On the worker side, we impose the additive structure

$$\alpha_{xyz} \equiv \alpha_{xy} + \alpha_{xz} + \alpha_{yz}. \tag{14}$$

This specification decomposes workers' preferences into three easily interpretable components. The term α_{xy} allows workers of different types to value firm types differently. The term α_{xz} captures heterogeneity across worker types in preferences over hours arrangements. Finally, α_{yz} allows the value of different hours arrangements to vary systematically across firm types.

On the firms' side, we analogously specify

$$\rho_{xyz} \equiv \rho_{xy} + \rho_{xz} + \rho_{yz}. \tag{15}$$

This specification similarly separates productivity into three components. The term ρ_{xy} captures systematic differences in productivity across firm types for workers of different types. The term ρ_{xz} allows the productivity value of different hours arrangements

[22] See Chapter 13 in Greene (2012) for a textbook discussion on minimum-distance estimation.

to vary across worker types. Finally, ρ_{yz} allows the productivity value of hours to vary across firm types.

The decompositions in Equations (14) and (15) are not separately identified without additional normalizations. More specifically, their individual components are identified only up to arbitrary additive shifts. For example, for any δ_x, the composite term α_{xyz} is unchanged if α_{xy} is replaced by $\alpha_{xy} + \delta_x$ and α_{xz} by $\alpha_{xz} - \delta_x$. The normalizations we impose in the empirical application are detailed in Appendix B.2.

4 Data

This section introduces our main data sources and describes the construction of worker and firm types. Additional details are provided in Appendix B.1.

Sample selection. The primary dataset is the latest three waves of the American Community Survey (ACS, 2022–2024), provided by the Integrated Public Use Microdata Series (IPUMS). We use the ACS to construct a sample of the high-end labor force in the United States as in Goldin (2014). The sample consists of individuals between 25 and 64 years old who hold at least a bachelor's degree. We distinguish between two subsamples. The employed subsample are those that work at least 35 hours per week and 40 weeks per year in a high earning occupation, as listed in Table 1. This corresponds to the full-time and full-year restriction, and the same set of occupations, that is also studied in Goldin (2014). Because most individuals in this segment of the labor market work full-time and full-year, this removes just 12.6% of the sample. We also drop 2.5% outliers of the earnings distribution on both ends. The unemployed subsample is also restricted to those that hold at least a bachelor's degree.

The second main dataset consists of individual job vacancies, and is provided by the National Labor Exchange (NLx) Research Hub. They collect online postings from companies' career websites, state job banks, and the federal jobs portal. The NLx attempts to cover all online vacancies in the United States, and represents roughly half of all openings as estimated in the Job Openings and Labor Turnover Survey (JOLTS) by the Bureau of Labor Statistics (BLS).[23] We use the NLx data to construct occupational vacancy distributions across regions, using almost 30 million unique postings. We then scale the distributions using adjusted JOLTS estimates on the number of vacancies per unemployed worker and the turnover rate.

[23] Comparable commercial vacancy data from the Burning Glass Institute with similar coverage has recently been used by Deming and Noray (2020) and Acemoglu, Autor, Hazell, and Restrepo (2022).

Table 1: Types of Occupations

Business	**Health**	**Tech**	**Science**	**Other**
Lawyers and Judges	Medical Practicioners	Comp. Programmers	Natural Scientists	Training & Teaching
Exec's and GMs	Adv. Clin. Prof.	Engineers	Social Scientists	Other
Managers	Clin. Prof.	Cybersec. & Infrastr.		
Financial Ops				
Business Ops				
Sales				

Notes. Classification of occupations into types. The individual occupations within each type can be found in Appendix B.1.

Worker and firm types, and hours. We group workers, firms, and hours into discrete categories based on observable characteristics. Workers are classified into $|\mathcal{X}| = 16$ types along four dimensions: gender, age (younger or older than 40), education (Bachelor's only or Master's and above, MA$^+$), and undergraduate field, which we split between *Science, Technology, Engineering, and Mathematics* (STEM) and *Arts, Humanities, and Social Sciences* (AHSS). Occupations are grouped into $|\mathcal{Y}| = 16$ types, following the classification of Goldin (2014) that distinguishes *Business*, *Health*, *Tech*, *Science*, and *Other* occupations. Table 1 lists the occupational types, and Appendix B.1 provides the underlying detailed occupations. Hours are classified as regular (35 to 49 per week) or long (50 or more), following Cortés and Pan (2019). Average hours within each category are comparable across occupations and worker types, at roughly 40 and 55 hours per week, respectively. Taken together, the model has $|\mathcal{X}| \times |\mathcal{Y}| \times |\mathcal{Z}| = 16 \times 16 \times 2 = 512$ submarkets.

Segmented labor markets. Following Ciscato and Goussé (2024), we define markets as the four U.S. macro-regions (Northeast, Midwest, South, and West). Variation in the composition of workers and firms across these markets identifies the search frictions, and yields substantial overidentification: the number of empirical moments exceeds the number of structural parameters. Appendix B.3 documents substantial variation in the marginal distribution of worker and firm types across macro-regions.

5 Estimates and Decompositions

We begin our discussion of the results by illustrating the in-sample and out-of-sample model fit and by presenting the main parameter estimates. We then decompose the long-hours premium, examining in each occupation how much of the earnings difference reflects firm-side productivity premiums, as opposed to compensation for workers' cost of longer hours or labor market tightness. Having established why long-hours premiums emerge, we turn to their implications for gender inequality.

5.1 Model Fit and Estimates

Model fit. Figure 2 illustrates the in-sample model fit by plotting the observed matching patterns and wages in each submarket against their predicted model counterparts. The diagonal line represents a perfect fit. Panel (a) shows that the model replicates matching patterns very well, with a correlation between observed and predicted values of 0.98; this also implies that hours in the model align closely with those in the data. Panel (b) shows that the model also replicates earnings well, with a correlation of 0.86.

Figure 2: Model Fit: Matches and Earnings

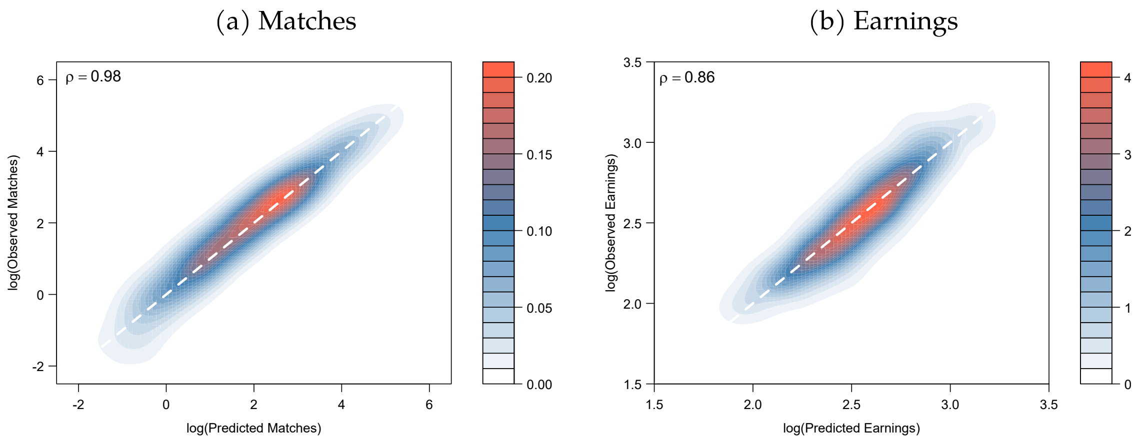

Notes. This figure shows the model predictions against the data moments. Panel (a) highlights the fit of matches, and Panel (b) the fit of earnings. Both matching patterns and earnings are expressed in logs. The color scale highlights the density of observations obtained from a kernel estimator.

In Figure 3, we document how closely our model matches non-targeted moments in the data. Panel (a) plots the model-implied log long-hours premium against its observed counterpart by occupation. The correlation is high, and the model reproduces the ranking across occupations well: occupations with large observed long-hours premiums in Figure 1, such as *Financial Operations*, *Lawyers and Judges*, and *Executives and General Managers*, also appear at the top of the model-implied premiums. Panel (b) shows the share of workers in each occupation working long hours, both in the data and as implied by the model. Reassuringly, the model again matches these shares closely.

We also validate the model-implied tightness. Although tightness is not directly observable in the data, we can construct rough proxies that are closely tied to market-level tightness. Specifically, we compute two measures: the match-to-vacancy ratio, defined as the number of matches M_{xyz} in submarket (x, y, z) divided by the number

Figure 3: Model Fit: Long-Hours Premium and Long Hours by Occupation

(a) Long-Hours Premium (b) Long Hours

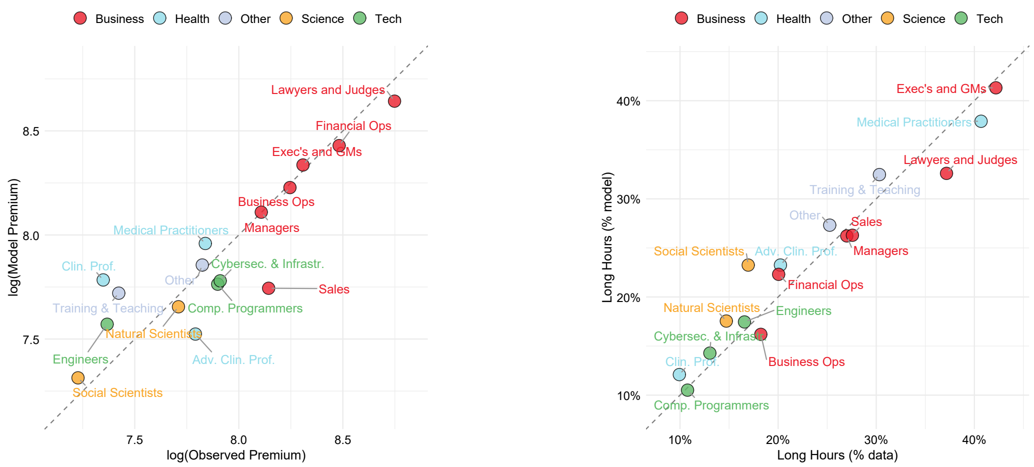

Notes. This figure shows the model predictions against the data moments. Panel (a) compares observed and predicted long-hours premiums by occupation, and Panel (b) compares the share of workers working long hours. The hours premium is expressed in logs.

of unfilled vacancies in occupation y; and the match-to-unemployment ratio, which instead divides by the number of unemployed workers of type x. Figure 4 presents the model-implied ratios. Both the match-to-vacancy and the match-to-unemployment ratios align closely with their data counterparts.

Figure 4: Model Fit: Match-to-Vacancy and Match-to-Unemployment Ratio

(a) Match-to-Vacancy by Occupation Type (b) Match-to-Unemployment by Worker Type

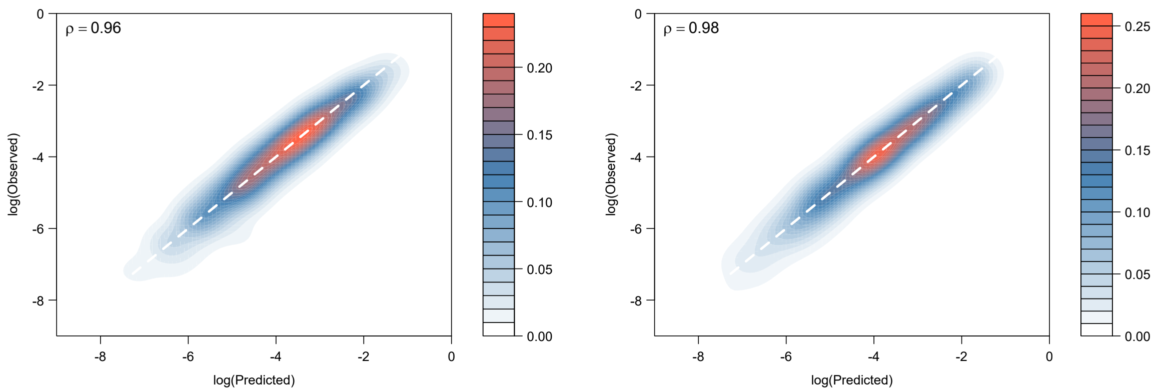

Notes. This figure shows the model predictions against the data moments. Panel (a) highlights the fit of the match-to-vacancy ratio by occupation, and Panel (b) the fit of the match-to-unemployment ratio by worker type. Both predicted and observed ratios are expressed in logs. The color scale highlights the density of observations obtained from a kernel estimator.

Productivity estimates. Figure 5 presents estimates of the productivity premium firms derive from employees working long hours, as measured by ρ_{yz}. Several occupations with large earnings premiums, such as *Lawyers and Judges* and *Executives and General Managers*, also exhibit large productivity premiums, consistent with the narrative in Goldin (2014) that occupational long-hours earnings premiums reflect production technologies. However, the relationship is far from one-to-one. The clearest deviation is in the *Tech* occupations: *Cybersecurity & Infrastructure* and *Computer Programmers* exhibit sizable long-hours premiums of close to 20% (see Figure 1), yet show little to no added productive value from working longer hours. As we show below, the bulk of this premium reflects tightness in long-hours markets rather than productivity gains. This deviation cautions against conflating productivity and earnings premiums and motivates the equilibrium decomposition that follows.

Figure 5: Estimates: Productivity Premium by Occupation

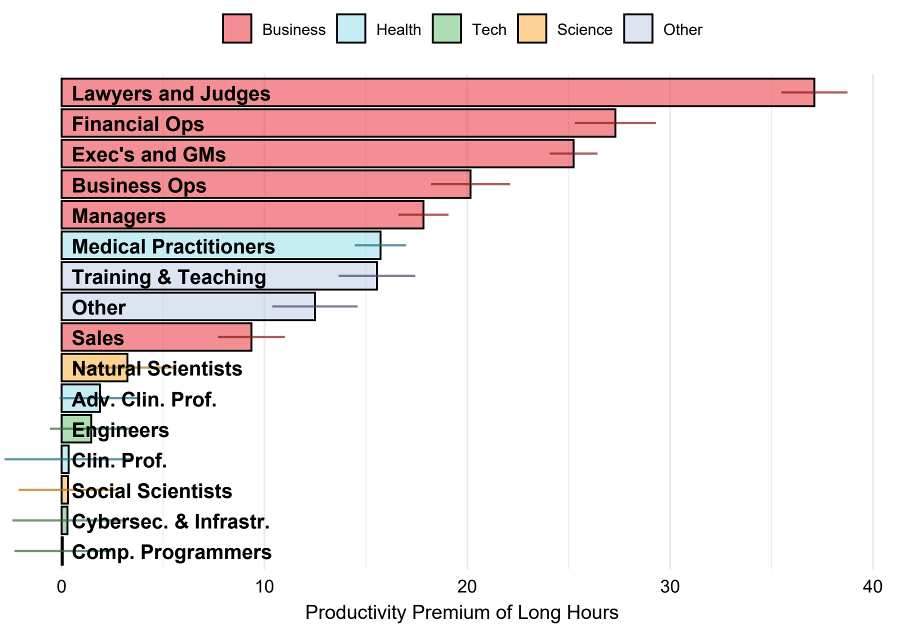

Notes. This figure shows the estimated productivity premium of long hours (ρ_{yz}) for each occupation, expressed as a percentage of regular-hour productivity. Higher values indicate occupations where long hours generate proportionally more output. The 95% confidence intervals are based on the asymptotic variance of the minimum distance estimator.

Amenity estimates. Figure 6 presents estimates of the utility workers derive from working regular hours by worker type, as measured by α_{xz}. Several patterns stand out. First, women value regular hours more than men do: conditional on occupation (that is, fixing α_{yz}), men are willing to give up about 19% of their wages to work regular rather than long hours, compared with about 26% for women. These estimates are consistent with prior findings that women place a higher value on shorter hours

Figure 6: Estimates: Amenity Premium by Worker Type

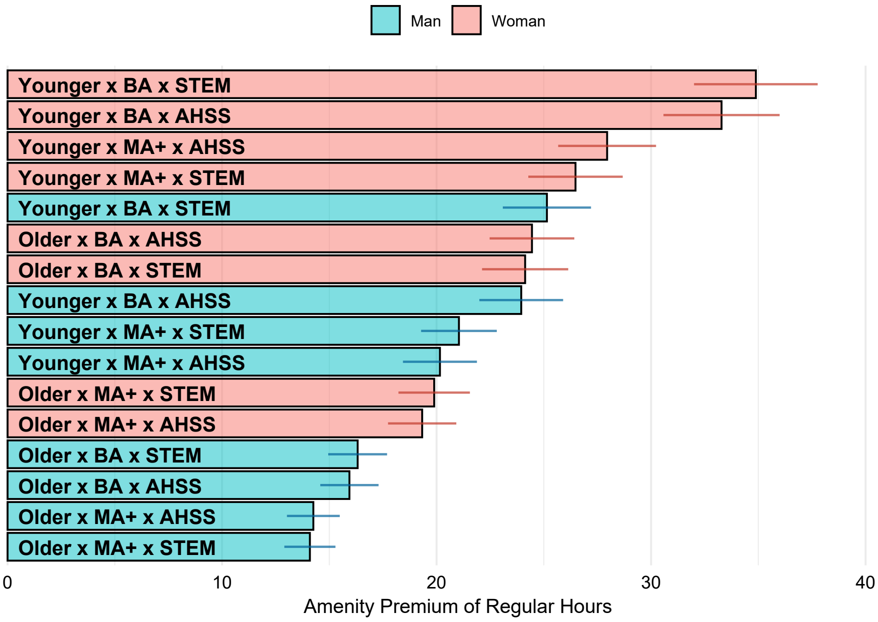

Notes. This figure shows the estimated utility cost of working long hours (α_{xz}) for each worker type, expressed as a percentage of mean wages in long-hours jobs. Higher values indicate a greater disutility from long hours relative to regular hours. The 95% confidence intervals are based on the asymptotic variance of the minimum distance estimator.

(Flabbi & Moro, 2012; Wiswall & Zafar, 2018), and younger women with a STEM BA exhibit the highest premium for regular hours. Second, conditional on gender, younger workers dislike long hours more than older workers do, likely reflecting constraints on flexibility associated with young children at home.[24] Third, conditional on gender and undergraduate field, workers with MA$^+$ degrees value regular hours less than those with only a BA. This may reflect selection—individuals anticipating a higher tolerance for long hours invest more in graduate education—or differences in the anticipated returns to long hours across degree levels. Finally, conditional on gender, age, and educational attainment, the differences between STEM and AHSS workers are relatively small, with the gap amounting to only a few percentage points in most cells. This limited heterogeneity in preferences over regular hours across undergraduate fields implies that most of the observed cross-field variation in long-hours shares (see Figure B.1 in the Appendix) reflects occupational sorting.

In Appendix B.3, we also present estimates for workers' preferences for regular hours by occupation, as measured by α_{yz}. We find that the disutility for working long

[24] This is consistent with the literature on labor supply changes around parenthood (Adda et al., 2017; Kleven et al., 2019). We cannot identify this channel directly, because our model does not incorporate family dynamics, though our estimates remain consistent with this interpretation. An alternative mechanism is a difference in preferences, namely that younger workers systematically place greater value on work-life balance than older workers.

hours is especially large in some *Business* occupations. However, contrasting these estimates with the corresponding productive value by occupation, as measured by ρ_{yz}, the former are smaller than the latter.

Elasticity of substitution estimate. The elasticity of substitution is estimated at $\lambda = -0.74$. To understand what this means in practice, consider its implications for how the match surplus is divided between workers and firms. When workers and firms are equally numerous, workers capture exactly 50% of the surplus. When workers outnumber firms two-to-one this falls to 37%, and when they outnumber firms four-to-one, to just 26%. Market tightness, in other words, plays a central role in determining bargaining outcomes.

External validation: Willingness-to-Pay. Given the central role of amenities, we provide external validation of these parameters. To that end, we compare our amenity estimates to the preference estimates from Maestas et al. (2023), who design a vignette experiment within the American Working Conditions Survey (AWCS). Respondents are presented with pairs of hypothetical job profiles that vary randomly across several attributes and asked to select their preferred option. We restrict the sample to respondents with at least a bachelor's degree and to choice pairs in which both jobs require at least 35 hours per week. Jobs offering 35–49 hours are classified as regular-hours jobs, and those offering 50 or more as long-hours jobs, matching our main definitions. We estimate a conditional logit model on the binary choice data with log earnings, a long-hours indicator, and a vector of other job amenities, separately for different worker types and sectors. Willingness to pay for regular hours is recovered as the compensating earnings differential implied by the estimated coefficients. Because the two methods produce estimates on different scales, we demean both within each dimension.[25] Figure 7 plots the demeaned estimates against each other. The raw correlation is 0.43: subgroups for which our model predicts a higher willingness to pay for regular hours, such as younger workers, women, and those in science occupations, also display stronger preferences for shorter hours in the survey experiment. Removing *Younger Men* from both samples raises the correlation substantially, to 0.77.

[25]The two approaches differ in two main ways. First, conceptually: our model captures total compensation differences between regular-hours and long-hours jobs, whereas the survey isolates the hours effect experimentally. Second, in terms of data and sample: our estimates rely on a selected sample of high-earning college graduates, while the AWCS asks for sectors and not occupations (at a coarse two digit level), and contains just 824 respondents (as few as 30 in the smallest cell).

Figure 7: External Validation: Willingness-to-Pay

Notes. This figure shows standardized WTP estimates for regular hours from our structural model against survey-based estimates from the AWCS (Maestas et al., 2023). Both sets of estimates are standardized within dimension (gender $\times$ age and occupation). The raw correlation is 0.43; removing *Younger Men* from both samples significantly increases the correlation to 0.77.

5.2 Decompositions

Long-hours earning premiums. To better understand how equilibrium long-hours premiums are shaped, we introduce a decomposition based on the equilibrium wage function in Equation (5). Let z denote a long-hours arrangement and z' a regular-hour arrangement. Let $\mu_x \equiv \frac{N_x}{\sum_x N_x}$ denote the relative frequency of type x in the population. The equilibrium long-hours premium across occupations (averaged across worker types) is defined as $\Delta_z W_y \equiv \sum_x \mu_x (W_{xyz} - W_{xyz'})$, which can be decomposed into contributions of three components:

$$\Delta_z W_y = \sum_x \mu_x \Bigg(\underbrace{(\rho_{xyz} - \rho_{xyz'})(1 - F(\theta_{xyz}))}_{\Delta \text{ Productivity}} + \underbrace{(\alpha_{xyz'} - \alpha_{xyz})F(\theta_{xyz})}_{\Delta \text{ Amenity}} + \underbrace{(F(\theta_{xyz'}) - F(\theta_{xyz}))\Phi_{xyz'}}_{\Delta \text{ Bargaining Power}} \Bigg). \tag{16}$$

See Appendix A.5 for the derivation.

The first component captures differences in firms' production technologies. The long-hours premium is larger the more productive long hours are relative to the regular-hour alternative ($\rho_{xyz} > \rho_{xyz'}$). How much of this additional productive value is passed on to workers, however, depends on firms' bargaining weights $F(\theta_{xyz})$. The second component captures the amenity loss from moving from regular to long hours. When long hours substantially reduce workers' non-material benefits ($\alpha_{xyz'} > \alpha_{xyz}$), the long-hours premium increases, again with the magnitude depending on bargaining power. The final term reflects how market tightness directly affects the division of surplus. When long-hours markets are tighter ($\theta_{xyz} > \theta_{xyz'}$), firms have less bargaining power than in regular-hour markets ($F(\theta_{xyz}) < F(\theta_{xyz'})$), and capture a smaller share of the joint surplus $\Phi_{xyz'}$.

Figure 8 presents the results from this decomposition. From the first column, it is apparent that productivity drives a substantial part of the long-hours premium in high-earning *Business* occupations: the productivity channel inflates equilibrium earnings by around 30 to 37%. In contrast, in other occupations such as *Tech*, productivity makes almost no contribution to the earnings premium. The second column shows that amenity values are instead a key driver of the long-hours premium: across all occupations, the disutility of working long hours is one of the largest contributing factors. This is notably the case in *Training & Teaching*, *Science*, and *Advanced Clinical Professionals*, where the long-hours premium simply reflects compensation for the additional costs workers incur when working long hours. The amenity contribution for *Advanced Clinical Professionals* exceeds 100%, which partly reflects the fact that this is one of the occupations for which our model overestimates the long-hours premium, as illustrated in Figure 3.[26]

The third column highlights the contribution of bargaining power. When long-hours markets are tighter than regular-hours markets, the long-hours premium is inflated. This is especially the case in occupations where we documented large excess demand, such as *Tech* occupations. Consistent with our productive-value estimates in Figure 5, these occupations exhibit almost no productive-value contribution to the earnings premium; yet because these markets are tight, workers are able to capture a large share of the surplus, as reflected in the substantial contribution of bargaining power (between 18 and 37%).

Gender earnings gap. We now decompose the gender earning gap. To that end, let x and x' be two types of workers that share the same characteristics except for gender:

[26] A plausible reason is that earnings in these occupations are driven more by institutional forces than by bargaining within our competitive search environment. Pay scales for *Advanced Clinical Professionals* are often set through collective bargaining or hospital and health-system salary structures, which compress the earnings variation that our model attributes to equilibrium responses to amenities and tightness. We retain the results for this occupation for completeness.

Figure 8: Decomposition: Long-Hours Earning Premium

	Productivity	Amenity	Bargaining Power	
Managers	37	60	3	Business
Sales	34	46	20	
Exec's and GMs	34	69	-3	
Financial Ops	32	70	-2	
Business Ops	31	67	2	
Lawyers and Judges	31	73	-4	
Medical Practitioners	40	61	-1	Health
Adv. Clin. Prof.	5	119	-24	
Clin. Prof.	1	72	27	
Engineers	5	77	18	Tech
Cybersec. & Infrastr.	1	62	37	
Comp. Programmers	0	63	37	
Natural Scientists	8	93	-1	Science
Social Scientists	2	50	48	
Training & Teaching	27	76	-2	Other
Other	21	86	-8	

Notes. This figure shows the decomposition of the long-hours premium across occupations (Equation (16)), expressed as a percentage of the total earnings difference. The first column shows the Δ Productivity term, the second column the Δ Amenity term, and the final column the Δ Bargaining Power term.

$x \equiv (m, \tilde{x})$ refers to a man with characteristics $\tilde{x}$ and $x' \equiv (w, \tilde{x})$ to a woman with those same characteristics. We define the average earnings gap as

$$\Delta G \equiv \sum_{y,z} M_{xyz} W_{xyz} - \sum_{y,z} M_{x'yz} W_{x'yz},$$

where wages are weighted by the number of workers across occupations and hours.

Following Oaxaca (1973) and Blinder (1973), we separate the wage gap into a sorting component and a wage-structure component. Building on the structural wage Equation (5), we then provide further insights into the different determinants of wages, as in Morchio and Moser (2026). Let $M_{xy} \equiv \sum_z M_{xyz}$ denote the total number of men of type x working in occupation y and similarly $M_{x'y} \equiv \sum_z M_{x'yz}$ the number of women of type x' working in occupation y. These marginals capture gender differences in sorting across occupations. Using them, we define the within-type shares across hours arrangements as $\sigma_{xyz} \equiv \frac{M_{xyz}}{M_{xy}}$ for men and $\sigma_{x'yz} \equiv \frac{M_{x'yz}}{M_{x'y}}$ for women, which capture differences in sorting across hours. We can then write

$$\Delta G = \sum_{y,z} \Bigg(\underbrace{(M_{xy} - M_{x'y})\sigma_{xyz}}_{\Delta \text{ Sorting in } y} + \underbrace{(\sigma_{xyz} - \sigma_{x'yz})M_{x'y}}_{\Delta \text{ Sorting in } z} \Bigg) W_{x'yz} +$$
$$\sum_{y,z} \Bigg(\underbrace{(\rho_{xyz} - \rho_{x'yz})(1 - F(\theta_{xyz}))}_{\Delta \text{ Productivity}} + \underbrace{(\alpha_{x'yz} - \alpha_{xyz})F(\theta_{xyz})}_{\Delta \text{ Amenity}} \tag{17}$$
$$+ \underbrace{(F(\theta_{x'yz}) - F(\theta_{xyz}))\Phi_{x'yz}}_{\Delta \text{ Bargaining Power}} \Bigg) M_{xyz}.$$

See Appendix A.5 for the derivation.

The results from this decomposition are presented in Table 2. The upper panel shows that the gender earnings gap in the data is about 21%, and the model predicts this gap almost perfectly, at about 20%. The lower panel shows how much of the earnings gap is explained by each component. Sorting into different occupations and hours is an important determinant of the gender earnings gap. If men and women worked in the same occupations, the gap would shrink by about 25%. The gender difference in sorting into long-hours positions also explains a sizable share, around 7%.[27] Both model-implied contributions are also close to the shares explained in the data.

Table 2: Decomposition: Gender Earnings Gap

	Data	**Model**
Earnings Gap	20.8%	20.4%
Decomposition		
Δ Sorting: Occupations	19.2%	24.5%
Δ Sorting: Hours	7.6%	6.5%
Earnings Structure	73.3%	
Δ Earnings: Productivity		47.1%
Δ Earnings: Amenities		21.5%
Δ Earnings: Bargaining Power		-2.5%
Remainder: Unexplained		2.7%

Notes. This table shows the gender gap in earnings. The first column applies a standard Oaxaca-Blinder decomposition, splitting the gap into differential sorting across occupations and hours, in addition to within-cell differences. The second column further decomposes the within-cell component using the structural wage Equation (5). All components are expressed as a percentage of the observed data gap.

[27]These magnitudes are broadly in line with the literature. Blau and Kahn (2017) and Olivetti et al. (2024) attribute roughly a third of the wage gap to occupational sorting, while Olivetti et al. (2024) find that sorting into long hours accounts for about 10% in the overall population.

Building on our wage equation, we can further decompose the earnings structure. First, the productivity channel makes a large contribution to the gender earnings gap, about 47%. These differences need not reflect actual productivity gaps by gender (e.g., from job tenure or job-specific human capital); they may equally capture statistical discrimination in how firms perceive men's and women's productivity, or taste-based discrimination.[28] Since our framework cannot separate these sources, the 47% captures their combined contribution. Second, a substantial gender gap in the cost of work, reflecting differences in amenities or constraints, explains about 22% of the earnings gap. Finally, bargaining power has a mildly dampening effect on wage inequality, in that women have, on average, slightly higher bargaining power.[29] In particular, women tend to sort into slightly tighter submarkets, thereby increasing their relative share of the surplus.

6 Counterfactual Policy Experiments

We now conduct two counterfactual policy experiments to examine how the gender earnings gap responds to interventions on opposite sides of the labor market. The first operates on the demand side, eliminating the gender gap in how firms evaluate workers' productive value. The second operates on the supply side, reallocating women from AHSS to STEM degrees.

Experiment 1: Closing the firm-side gender gap. The gender earnings decomposition above showed that about 47% of the gap can be explained by firms' evaluation of productive value (see Table 2). Given recent evidence from Sin et al. (2022), who find only a modest contribution of gender gradients in productivity, most of these differences in firms' valuations might reflect statistical or taste-based discrimination rather than actual productivity gaps. From a policy perspective, it therefore seems preferable to target this part of the gap. Because the share attributable to discrimination is uncertain, we close this gender wedge to varying degrees in the exercise that follows.

More formally, define $x' \equiv (w, \tilde{x})$ and $x \equiv (m, \tilde{x})$ as above; that is, x is a male worker with characteristics $\tilde{x}$ and x' is a female worker with the same characteristics. We adjust the productive value for women of type x' in occupation y working z hours as follows:

$$\tilde{\rho}_{x'yz} \equiv (1-\tau)\rho_{x'yz} + \tau\rho_{xyz},$$

[28] One form of statistical discrimination stems from missing out on promotions around the typical ages of first childbirth, which lowers the likelihood of within-firm promotions even for women who never have children (Bronson & Skogman Thoursie, 2019).

[29] As noted earlier, these differences do not reflect gender differences in wage negotiation, which have been identified as an alternative channel for gender wage gaps (Bertrand, 2011; Leibbrandt & List, 2015), but are instead market-level differences stemming from how men and women sort across submarkets.

where $\tau \in [0, 1]$ captures the degree of convergence. When $\tau = 1$, the productive values of women and men are equalized, conditional on the same characteristics $\tilde{x}$, occupation y, and hours z.

The first column of Table 3 reports the baseline gender earnings decomposition, replicating Table 2. Moving from the second to the sixth column, we gradually increase the extent to which the gap in productive values is closed. Fully closing this gap eliminates 69% of the baseline gender gap. Several other patterns stand out. First, the productivity component declines smoothly with the convergence parameter, falling mechanically to zero once the entire productive value is equalized. Second, occupational sorting falls steadily as women re-sort into occupations that previously carried a (larger) productivity wedge. The amenity channel, in turn, becomes somewhat more important, since closing the productivity channel leads women to sort relatively more on the disutility of long hours. Third, because the perceived productivity gap between men and women shrinks, firms' relative demand for women rises, and this heightened competition for female workers across occupations raises labor-market tightness. Fully closing the gap, this bargaining channel further narrows the gender earnings gap by about 15%.

Table 3: Decomposition: Counterfactual Gender Earnings Gap

Convergence rate τ	$\tau = 0.0$	$\tau = 0.2$	$\tau = 0.4$	$\tau = 0.6$	$\tau = 0.8$	$\tau = 1.0$
Gender gap	20.38	17.85	15.19	12.42	9.53	6.53
Contributions:						
Δ Sorting: Occupations	24.5 %	24.3 %	23.2 %	21.6 %	19.3 %	16.4 %
Δ Sorting: Hours	6.5 %	6.6 %	6.6 %	6.6 %	6.6 %	6.6 %
Δ Earnings: Productivity	47.1%	37.5 %	27.8 %	18.3 %	8.9 %	0.0 %
Δ Earnings: Amenities	21.5 %	21.9 %	22.2 %	22.4 %	22.6 %	22.9 %
Δ Earnings: Bargaining Power	-2.5 %	-4.5 %	-6.7 %	-9.1 %	-11.8 %	-14.6 %
Share of baseline gap eliminated	0.00	12.93	26.32	40.06	54.18	68.68

Notes. This table shows the counterfactual gender earnings gap for varying convergence rates τ. Each column decomposes the within-cell component using the structural wage equation. All components are expressed as a percentage of the observed data gap.

A significant general-equilibrium effect runs through market tightness: by tightening the markets where women operate and loosening those for men, the policy experiment strengthens women's bargaining power while weakening men's on average. Figure B.9 in the Appendix displays the re-sorting patterns by gender across occupations. Women move into occupations with higher baseline tightness, such as *Tech*. Men, by contrast, face stiffer competition, since firms now treat men and women as more substitutable than in the baseline economy, and consequently re-sort into occupations with more slack (e.g., *Training & Teaching*, *Advanced Clinical Professionals*, and *Natural Scientists*). Consistent with this, Figure B.10 in the Appendix shows that tightness rises for women in almost all occupations and falls for men in nearly all of them.

Experiment 2: Closing the gender gap in STEM. Our second policy experiment operates on the supply side, raising the share of women in STEM by shifting them from AHSS to STEM fields while holding workers' preferences and firms' assessment of productive value constant. Specifically, we move 20% of women in each age-by-education cell from AHSS to STEM.[30] Relative to the ACS composition, this amounts to closing about 70% of the gender gap in STEM fields, a substantial narrowing of the gap.

Figure 9 decomposes the change in average earnings into within-field and compositional components. For women, the overall effect of the policy is to raise average earnings by about 0.7%.[31] This change decomposes into within-field contributions of about -0.03% from AHSS women and -1.06% from STEM women, alongside a positive compositional contribution, so the net increase is driven almost entirely by compositional effects.[32] The reason is that, on average, women in STEM sort into relatively higher-earning occupations, and this compositional gain dominates even though the larger supply of STEM women lowers average earnings within STEM. For the effects of average earnings of women at the occupation level, see Figure B.11

For men, average earnings fall slightly, by about 0.41%, with AHSS and STEM men each contributing around -0.2%. The driver is the equilibrium rebalancing of tightness across submarkets: STEM-intensive occupations become more crowded as women with STEM degrees enter, intensifying competition for STEM men. This competition spills over to AHSS men as well, particularly in *Business* occupations, where AHSS men's baseline employment shares are largest. This pattern is illustrated in Figure B.13 in the Appendix. Overall, the gender earnings gap drops by about 1 percentage point.

Discussion and comparison. These experiments yield two main takeaways. First, general-equilibrium effects are quantitatively important, and their sign is ambiguous a priori: depending on the intervention, they can either amplify or attenuate the direct impact. In the first policy, the general-equilibrium response reinforced the direct effect, adding a further 15% of gap closure through the bargaining channel. In the second, it worked in the opposite direction. While shifting women into STEM does narrow

[30]More formally, let $x \equiv (w, \tilde{x}, \text{AHSS})$ denote a woman with characteristics $\tilde{x}$ (age and educational attainment) and an AHSS major, and let $x' \equiv (w, \tilde{x}, \text{STEM})$ denote the otherwise identical woman with a STEM major. The new numbers of women of each type are then $\tilde{N}_x = 0.8N_x$ and $\tilde{N}_{x'} = N_{x'} + 0.2N_x$ for all $\tilde{x}$, reflecting a shift of 20% of AHSS women into STEM.

[31]This is much smaller than the first policy's effect: that policy acts on the perceived productive-value gap and the associated tightness adjustments, whereas this one shifts only composition. A caveat: moving women into STEM may also have indirect benefits through role-model effects.

[32]To be more specific, we calculate this compositional effect by averaging out the changes in the share of women in each undergraduate field, multiplied by the average earnings for women in that field.

Figure 9: Decomposition: Counterfactual Earnings Difference

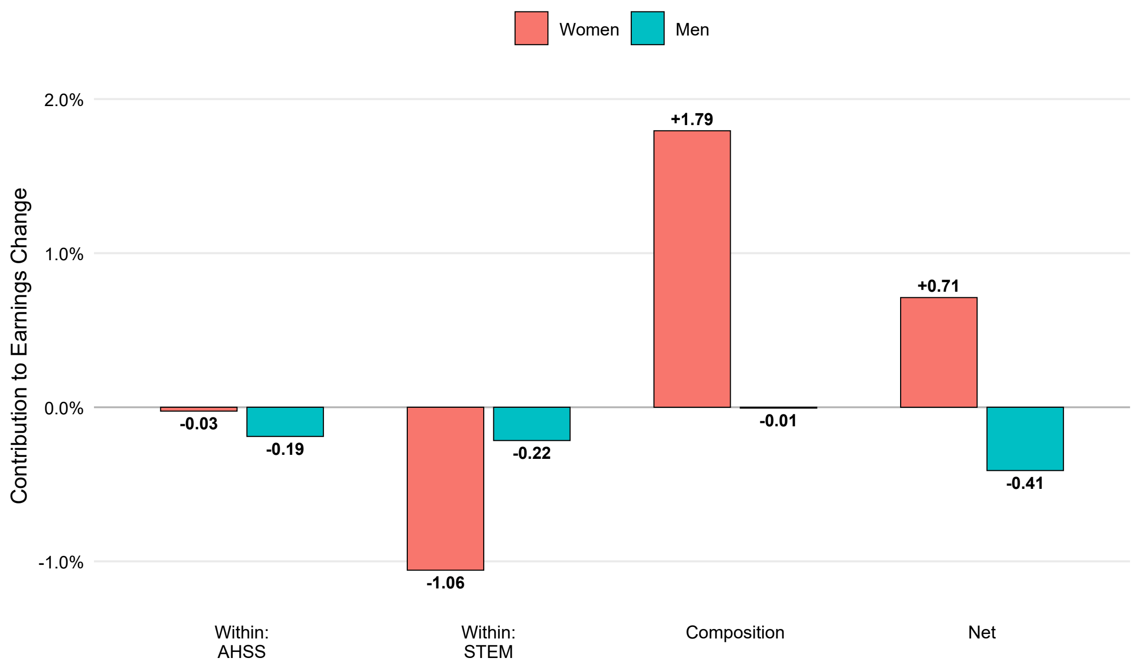

Notes. This figure shows the difference in earnings between the baseline and counterfactual equilibria. The first and second columns highlight how much of the total earnings difference is due to earnings changing for individuals with an AHSS degree. Similarly, the third and fourth column show the contribution of within-STEM earnings differences. The fifth and sixth column capture the contribution of the composition shifting from AHSS toward STEM degrees. The final two columns show the total difference in earnings.

the gap in occupational sorting, the resulting inflow congests the submarkets these women enter, loosening them and eroding workers' bargaining power, which partially offsets the direct gains to women's earnings. The same congestion mechanism that strengthens women's position in the first experiment thus weakens it in the second.

Second, the gender earnings gap is fundamentally a pricing problem rather than a sorting problem. The decomposition in Table 2 attributes a larger share of the gap to perceived gender differences in productive value than to sorting across occupations. Reallocating women across fields therefore raises their average earnings through better placement, but cannot close the residual gap with, for instance, STEM men, which stems from how their productive value is priced rather than from where they are employed. Closing that gap requires acting on firms' valuations directly.

7 Conclusion

This paper studies why long-hours premiums arise and why they vary across occupations. Previous work emphasizes firm-side productivity premiums (Goldin, 2014);

we instead highlight the role of workers' amenities and of the environment in which workers and firms interact. We develop a tractable equilibrium search-and-matching framework that combines directed search, as in Moen (1997), with two-sided unobserved heterogeneity, as in Choo and Siow (2006). The model delivers an intuitive equilibrium wage function that decomposes the long-hours premium into its components. Two counterfactual policy experiments underscore the importance of accounting for general-equilibrium effects.

An important direction for future research is to further unpack the determinants of workers' amenities and firms' productivity. On the worker side, the value placed on long hours may reflect genuine preferences, prevailing societal norms, and binding constraints such as childcare responsibilities. On the firm side, measured productivity values may similarly conflate true productivity differences, employers' perceptions of productivity, and taste-based discrimination. Separating these channels would further sharpen the welfare and policy analysis, since the appropriate response may depend on which one dominates. A further valuable direction would extend the model to incorporate human capital accumulation and dynamic returns to long hours, capturing how today's hours shape future earnings.

References

Acemoglu, D., Autor, D., Hazell, J., & Restrepo, P. (2022). Artificial intelligence and jobs: Evidence from online vacancies. *Journal of Labor Economics*, *40*(S1), S293–S340.

Adda, J., Dustmann, C., & Stevens, K. (2017). The career costs of children. *Journal of Political Economy*, *125*(2), 293–337.

Aguiar, M., & Hurst, E. (2007). Measuring trends in leisure: The allocation of time over five decades. *The Quarterly Journal of Economics*, *122*(3), 969-1006.

Berry, S., Gandhi, A., & Haile, P. (2013). Connected substitutes and invertibility of demand. *Econometrica*, *81*(5), 2087–2111.

Bertrand, M. (2011). New perspectives on gender. In D. Card & O. Ashenfelter (Eds.), (Vol. 4, p. 1543-1590). Elsevier.

Bertrand, M., Goldin, C., & Katz, L. F. (2010). Dynamics of the gender gap for young professionals in the financial and corporate sectors. *American Economic Journal: Applied Economics*, *2*(3), 228–255.

Bick, A., Blandin, A., & Rogerson, R. (2022). Hours and wages. *The Quarterly Journal of Economics*, *137*(3), 1901–1962.

Black, D. A., Haviland, A. M., Sanders, S. G., & Taylor, L. J. (2008). Gender wage disparities among the highly educated. *Journal of Human Resources*, *43*(3), 630–659.

Blau, F. D., & Kahn, L. M. (2017). The gender wage gap: Extent, trends, and explanations. *Journal of Economic Literature*, *55*(3), 789–865.

Blinder, A. S. (1973). Wage discrimination: reduced form and structural estimates. *Journal of Human Resources*, 436–455.

Breda, T., Grenet, J., Monnet, M., & van Effenterre, C. (2023). How effective are female role models in steering girls towards STEM? *The Economic Journal*, *133*(653), 1773–1795.

Bronson, M. A., & Skogman Thoursie, P. (2019). *The wage growth and within-firm mobility of men and women: New evidence and theory.* (Mimeo)

Bureau of Labor Statistics. (2024). *Fastest growing occupations.* Retrieved from https://www.bls.gov/ooh/fastest-growing.htm (Accessed: 2025-04-09)

Caldwell, S., & Danieli, O. (2024). Outside options in the labour market. *Review of Economic Studies*, *91*(6), 3286–3315.

Card, D., Cardoso, A. R., & Kline, P. (2016). Bargaining, sorting, and the gender wage gap: Quantifying the impact of firms on the relative pay of women. *The Quarterly Journal of Economics*, *131*(2), 633-686.

Carrell, S. E., Page, M. E., & West, J. E. (2010). Sex and science: How professor gender perpetuates the gender gap. *The Quarterly Journal of Economics*, *125*(3), 1101–1144.

Chiappori, P.-A., Salanié, B., & Weiss, Y. (2017). Partner choice, investment in children, and the marital college premium. *American Economic Review*, *107*(8), 2109–2167.

Choo, E., & Siow, A. (2006). Who marries whom and why. *Journal of Political Economy*, *114*(1), 175–201.

Ciscato, E., & Goussé, M. (2024). *Matching on gender and sexual orientation* (IZA Discussion Paper No. 17420). Bonn: Institute of Labor Economics (IZA).

Cortés, P., & Pan, J. (2019). When time binds: Substitutes for household production, returns to working long hours, and the skilled gender wage gap. *Journal of Labor Economics*, *37*(2), 351–398.

Deming, D. J., & Noray, K. (2020). Earnings dynamics, changing job skills, and STEM careers. *The Quarterly Journal of Economics*, *135*(4), 1965–2005.

Dupuy, A., & Galichon, A. (2014). Personality traits and the marriage market. *Journal of Political Economy*, *122*(6), 1271–1319.

Dupuy, A., & Galichon, A. (2022). A note on the estimation of job amenities and labor productivity. *Quantitative Economics*, *13*(1), 153–177.

Eeckhout, J., & Kircher, P. (2010). Sorting and decentralized price competition. *Econometrica*, *78*(2), 539–574.

Eeckhout, J., & Kircher, P. (2018). Assortative matching with large firms. *Econometrica*, *86*(1), 85–132.

Erosa, A., Fuster, L., Kambourov, G., & Rogerson, R. (2022). Hours, occupations, and gender differences in labor market outcomes. *American Economic Journal: Macroeconomics*, *14*(3), 543–590.

Erosa, A., Fuster, L., Kambourov, G., & Rogerson, R. (2024). Wage and earnings inequality between and within occupations: The role of labor supply. *Journal of Monetary Economics*, 103553.

Flabbi, L., & Moro, A. (2012). The effect of job flexibility on female labor market outcomes: Estimates from a search and bargaining model. *Journal of Econometrics*, *168*(1), 81–95.

Fox, J. T., Yang, C., & Hsu, D. H. (2018). Unobserved heterogeneity in matching games. *Journal of Political Economy*, *126*(4), 1339–1373.

Goldin, C. (2014). A grand gender convergence: Its last chapter. *American Economic Review*, *104*(4), 1091–1119.

Goldin, C., & Katz, L. F. (2016). A most egalitarian profession: pharmacy and the evolution of a family-friendly occupation. *Journal of Labor Economics*, *34*(3), 705–746.

Greene, W. H. (2012). *Econometric analysis* (7th ed.). Boston: Pearson Education.

Hershbein, B., & Kahn, L. B. (2018). Do recessions accelerate routine-biased technological change? evidence from vacancy postings. *American Economic Review*, *108*(7), 1737–1772.

Hosios, A. J. (1990). On the efficiency of matching and related models of search and unemployment. *The Review of Economic Studies*, *57*(2), 279–298.

Hwang, H.-s., Mortensen, D. T., & Reed, W. R. (1998). Hedonic wages and labor market search. *Journal of Labor Economics*, *16*(4), 815–847.

Jang, Y., & Yum, M. (2022). Nonlinear occupations and female labor supply over time. *Review of Economic Dynamics*, *46*, 51–73.

Kleven, H., Landais, C., & Søgaard, J. E. (2019). Children and gender inequality: Evidence from denmark. *American Economic Journal: Applied Economics*, *11*(4), 181–209.

Lamadon, T., Lise, J., Meghir, C., & Robin, J.-M. (2024). *Labor market matching, wages, and amenities* (Tech. Rep.). National Bureau of Economic Research.

Lang, K., & Majumdar, S. (2004). The pricing of job characteristics when markets do not clear: theory and policy implications. *International Economic Review*, *45*(4), 1111–1128.

Le Barbanchon, T., Rathelot, R., & Roulet, A. (2021). Gender differences in job search: Trading off commute against wage. *The Quarterly Journal of Economics*, *136*(1), 381-426.

Leibbrandt, A., & List, J. A. (2015). Do women avoid salary negotiations? Evidence from a large-scale natural field experiment. *Management Science*, *61*(9), 2016–2024.

Lewbel, A. (2019). The identification zoo: Meanings of identification in econometrics. *Journal of Economic Literature*, *57*(4), 835–903.

Maestas, N., Mullen, K. J., Powell, D., Von Wachter, T., & Wenger, J. B. (2023). The value of working conditions in the united states and implications for the structure of wages. *American Economic Review*, *113*(7), 2007–2047.

Mantovani, C. (2023). *Hours-biased technological change* (Tech. Rep.). Working Paper.

McFadden, D. (1978). Modelling the choice of residential location. *Transportation Research Record*, *672*, 72-77.

Menzio, G., & Shi, S. (2010). Block recursive equilibria for stochastic models of search on the job. *Journal of Economic Theory*, *145*(4), 1453–1494.

Moen, E. R. (1997). Competitive search equilibrium. *Journal of Political Economy*, *105*(2), 385–411.

Montgomery, J. D. (1991). Equilibrium wage dispersion and interindustry wage differentials. *The Quarterly Journal of Economics*, *106*(1), 163–179.

Morchio, I., & Moser, C. (2026). The gender pay gap: Micro sources and macro consequences. *American Economic Review*, *116*(5), 1765–1810.

Oaxaca, R. (1973). Male-female wage differentials in urban labor markets. *International Economic Review*, 693–709.

Olivetti, C., Pan, J., & Petrongolo, B. (2024). The evolution of gender in the labor market. In *Handbook of labor economics* (Vol. 5, pp. 619–677). Elsevier.

Peters, M. (1984). Bertrand equilibrium with capacity constraints and restricted mobility. *Econometrica*, 1117–1127.

Rosen, S. (1986). The theory of equalizing differences. *Handbook of labor economics*, *1*, 641–692.

Roth, A. E., & Sotomayor, M. (1989). The college admissions problem revisited. *Econometrica*, 559–570.

Rothenberg, T. J. (1971). Identification in parametric models. *Econometrica*, *39*(3), 577–591.

Ruggles, S., Flood, S., Sobek, M., Backman, D., Cooper, G., Rivera Drew, J., … Williams, K. (2025). *IPUMS USA: Version 16.0.* IPUMS.

Sin, I., Stillman, S., & Fabling, R. (2022). What drives the gender wage gap? examining the roles of sorting, productivity differences, bargaining, and discrimination. *The Review of Economics and Statistics*, *104*(4), 636-651.

Sloane, C. M., Hurst, E. G., & Black, D. A. (2021). College majors, occupations, and the gender wage gap. *Journal of Economic Perspectives*, *35*(4), 223–48.

Sorkin, I. (2018). Ranking firms using revealed preference. *The Quarterly Journal of Economics*, *133*(3), 1331–1393.

Sullivan, P., & To, T. (2014). Search and nonwage job characteristics. *Journal of Human Resources*, *49*(2), 472–507.

Taber, C., & Vejlin, R. (2020). Estimation of a Roy/search/compensating differential model of the labor market. *Econometrica*, *88*(3), 1031–1069.

Tamer, E. (2003). Incomplete simultaneous discrete response model with multiple equilibria. *The Review of Economic Studies*, *70*(1), 147-165.

Vattuone, G. (2026). *Worker sorting and the gender wage gap* (Tech. Rep.). Working Paper.

Wiswall, M., & Zafar, B. (2018). Preference for the workplace, investment in human capital, and gender. *The Quarterly Journal of Economics*, *133*(1), 457–507.

Wright, R., Kircher, P., Julien, B., & Guerrieri, V. (2021). Directed search and competitive search equilibrium: A guided tour. *Journal of Economic Literature*, *59*(1), 90–148.

Xiao, P. (2024). Equilibrium sorting and the gender wage gap. *Unpublished Manuscript*.

The Price of Long Hours

Online Appendix

A Theory

A.1 Derivation of Equation (5)

We now derive equilibrium wages in Equation (5). To that end, we first substitute out W from the participation constraint in Equation (4). This yields the following unconstrained optimization problem:

$$\max_{\theta_{xyz}} Q\left(\theta_{xyz}\right)\left(\Phi_{xyz}-\frac{U_x}{P\left(\theta_{xyz}\right)}\right),$$

where $\Phi_{xyz} \equiv \alpha_{xyz}+\rho_{xyz}$ denotes the total surplus in the match. Deriving the first-order condition (FOC) gives

$$\frac{\partial Q\left(\theta_{xyz}\right)}{\partial \theta_{xyz}}\left(\Phi_{xyz}-\frac{U_x}{P\left(\theta_{xyz}\right)}\right)+Q\left(\theta_{xyz}\right)\left(\frac{U_x}{P\left(\theta_{xyz}\right)^2} \frac{\partial P\left(\theta_{xyz}\right)}{\partial \theta_{xyz}}\right)=0.$$

Using the fact that $\frac{\partial Q(\theta_{xyz})}{\partial \theta_{xyz}}=\frac{P(\theta_{xyz})}{\theta_{xyz}}\left(\frac{1}{P(\theta_{xyz})} \frac{\partial P(\theta_{xyz})}{\partial \theta_{xyz}}-\frac{1}{\theta_{xyz}}\right)$, and multiplying both sides of the FOC by θ_{xyz}, we obtain

$$\left(1-\theta_{xyz} \frac{1}{P\left(\theta_{xyz}\right)} \frac{\partial P(\theta_{xyz})}{\partial \theta_{xyz}}\right)\left(\Phi_{xyz}-\frac{U_x}{P\left(\theta_{xyz}\right)}\right)-\theta_{xyz} \frac{U_x}{P\left(\theta_{xyz}\right)^2} \frac{\partial P(\theta_{xyz})}{\partial \theta_{xyz}}=0.$$

Substituting workers' market utility and using the definition of the job finding elasticity then yields Equation (5).

A.2 Auxiliary Results

Lemma A.1. *For all* (x, y, z) *and* $\boldsymbol{\theta} \in \Theta$,

$$\frac{\partial \overline{U}_{xyz}(\theta_{xyz})}{\partial \theta_{xyz}}>0, \quad \text{and} \quad \frac{\partial \overline{\Pi}_{xyz}(\theta_{xyz})}{\partial \theta_{xyz}}<0.$$

Proof. Since $\lambda<0$, we have that

$$\begin{aligned}
\frac{\partial \overline{U}_{xyz}(\theta_{xyz})}{\partial \theta_{xyz}} &= \Phi_{xyz} \frac{\partial}{\partial \theta_{xyz}}\left(P(\theta_{xyz})(1-F(\theta_{xyz}))\right) \\
&= \underbrace{\Phi_{xyz}}_{>0} \underbrace{\frac{\partial}{\partial \theta_{xyz}}\left(\left(1+\theta_{xyz}^{\lambda}\right)^{\frac{1-\lambda}{\lambda}}\right)}_{>0} \\
&> 0,
\end{aligned}$$

and

$$
\begin{aligned}
\frac{\partial \overline{\Pi}_{xyz}(\theta_{xyz})}{\partial \theta_{xyz}} &= \Phi_{xyz} \frac{\partial}{\partial \theta_{xyz}} (Q(\theta_{xyz}) F(\theta_{xyz})) \\
&= \underbrace{\Phi_{xyz}}_{>0} \underbrace{\frac{\partial}{\partial \theta_{xyz}} \left(\left(\frac{\theta_{xyz}^{\lambda}}{1+\theta_{xyz}^{\lambda}} \right)^{\frac{\lambda-1}{\lambda}} \right)}_{<0} \\
&< 0.
\end{aligned}
$$

□

Lemma A.2. *For all (x, y, z) and $\boldsymbol{\theta} \in \Theta$,*

$$
\frac{\partial \log S^{w}_{yz|x}(\boldsymbol{\theta})}{\partial \theta_{xyz}} > 0, \quad \textit{and} \quad \frac{\partial \log S^{f}_{xz|y}(\boldsymbol{\theta})}{\partial \theta_{xyz}} < 0,
$$

and for all $(x', y', z') \neq (x, y, z)$,

$$
\frac{\partial \log S^{w}_{yz|x}(\boldsymbol{\theta})}{\partial \theta_{x'y'z'}} \begin{cases} < 0, & \textit{if } x = x', \\ = 0, & \textit{otherwise}, \end{cases} \quad \textit{and} \quad \frac{\partial \log S^{f}_{xz|y}(\boldsymbol{\theta})}{\partial \theta_{x'y'z'}} \begin{cases} > 0, & \textit{if } y = y', \\ = 0, & \textit{otherwise}. \end{cases}
$$

Proof. By Lemma A.1,

$$
\begin{aligned}
\frac{\partial \log S^{w}_{yz|x}(\boldsymbol{\theta})}{\partial \theta_{xyz}} &= \underbrace{\left(1 - \pi^{w}_{yz|x}(\boldsymbol{\theta})\right)}_{>0} \underbrace{\frac{\partial \overline{U}_{xyz}(\theta_{xyz})}{\partial \theta_{xyz}}}_{>0} \\
&> 0,
\end{aligned}
$$

and

$$
\begin{aligned}
\frac{\partial \log S^{f}_{xz|y}(\boldsymbol{\theta})}{\partial \theta_{xyz}} &= \underbrace{\left(1 - \pi^{f}_{xz|y}(\boldsymbol{\theta})\right)}_{>0} \underbrace{\frac{\partial \overline{\Pi}_{xyz}(\theta_{xyz})}{\partial \theta_{xyz}}}_{<0} \\
&< 0.
\end{aligned}
$$

Similarly,

$$
\begin{aligned}
\frac{\partial \log S^{w}_{yz|x}(\boldsymbol{\theta})}{\partial \theta_{x'y'z'}} &= -\mathbb{I}[x' = x] \underbrace{\pi^{w}_{y'z'|x}(\boldsymbol{\theta})}_{>0} \underbrace{\frac{\partial \overline{U}_{xy'z'}(\theta_{xyz})}{\partial \theta_{x,y',z'}}}_{>0} \\
&\begin{cases} < 0, & \text{if } x = x', \\ = 0, & \text{otherwise}, \end{cases}
\end{aligned}
$$

and

$$\frac{\partial \log S^f_{xz|y}(\boldsymbol{\theta})}{\partial \theta_{x'y'z'}} = -\mathbb{I}[y' = y] \underbrace{\pi^f_{x',z'|y}(\boldsymbol{\theta})}_{>0} \underbrace{\frac{\partial \overline{\Pi}_{x'yz'}(\theta_{xyz})}{\partial \theta_{x',y,z'}}}_{<0}$$

$$\begin{cases} > 0, & \text{if } y = y', \\ = 0, & \text{otherwise.} \end{cases}$$

□

A.3 Proof of Proposition 1

We prove existence of a directed search equilibrium by means of Brouwer's fixed point theorem. We stack the system in Equation (9) across submarkets. The problem reduces to finding a fixed point $\boldsymbol{\theta} = \mathbf{T}(\boldsymbol{\theta})$, where the vector-valued map $\mathbf{T}$ is defined as

$$\mathbf{T} : \Theta \to \Theta : \boldsymbol{\theta} \mapsto \mathbf{T}(\boldsymbol{\theta}) = \frac{S^f_{xz|y}(\boldsymbol{\theta})}{S^w_{yz|x}(\boldsymbol{\theta})}.$$

Note first that both payoff terms $\overline{U}_{xyz}$ and $\overline{\Pi}_{xyz}$ are bounded for all submarkets:

$$\begin{aligned} \lim_{\theta_{xyz}\to 0} \overline{U}_{xyz}(\boldsymbol{\theta}) &= \Phi_{xyz} \lim_{\theta_{xyz}\to 0} P(\theta_{xyz})(1 - F(\theta_{xyz})) \\ &= \Phi_{xyz} \lim_{\theta_{xyz}\to 0} \left(1 + \theta^{\lambda}_{xyz}\right)^{\frac{1}{\lambda}} \frac{1}{(1 + \theta^{\lambda}_{xyz})} \\ &= 0, \end{aligned}$$

and

$$\lim_{\theta_{xyz}\to +\infty} \overline{U}_{xyz}(\boldsymbol{\theta}) = \Phi_{xyz}.$$

Similarly,

$$\begin{aligned} \lim_{\theta_{xyz}\to 0} \overline{\Pi}_{xyz}(\boldsymbol{\theta}) &= \Phi_{xyz} \lim_{\theta_{xyz}\to 0} Q(\theta_{xyz})F(\theta_{xyz}) \\ &= \Phi_{xyz} \lim_{\theta_{xyz}\to 0} \left(1 + \theta^{\lambda}\right)^{\frac{1-\lambda}{\lambda}} \theta^{\lambda-1} \\ &= \Phi_{xyz}, \end{aligned}$$

and

$$\lim_{\theta_{xyz}\to +\infty} \overline{\Pi}_{xyz}(\boldsymbol{\theta}) = 0.$$

Given the monotonicity result in Lemma A.1, we therefore obtain that $\overline{U}_{xyz}, \overline{\Pi}_{xyz} \in [0, \Phi_{xyz}]$ for all (x, y, z).

For all y, define $\overline{m}_y \equiv \max_{x',z'} \exp(\Phi_{x'yz'})$. Observe that $S^f_{xz|y}(\boldsymbol{\theta})$ is bounded as

$$\frac{N_y}{|\mathcal{X}||\mathcal{Z}|\overline{m}_y} \leqslant S^f_{xz|y}(\boldsymbol{\theta}) \leqslant \frac{N_y \overline{m}_y}{|\mathcal{X}||\mathcal{Z}|}. \tag{A.1}$$

Similarly, for all x, define $\overline{m}_x \equiv \max_{y',z'} \exp\left(\Phi_{xy'z'}\right)$. $S^w_{yz|x}\left(\boldsymbol{\theta}\right)$ is bounded as

$$\frac{N_x}{|\mathcal{Y}||\mathcal{Z}|\overline{m}_x} \leqslant S^w_{yz|x}\left(\boldsymbol{\theta}\right) \leqslant \frac{N_x \overline{m}_x}{|\mathcal{Y}||\mathcal{Z}|}. \tag{A.2}$$

Together, Equations (A.1) and (A.2) imply that, for all (x, y, z),

$$\underline{\theta}_{xyz} \leqslant \frac{S^f_{xz|y}\left(\boldsymbol{\theta}\right)}{S^w_{yz|x}\left(\boldsymbol{\theta}\right)} \leqslant \overline{\theta}_{xyz},$$

where $\underline{\theta}_{xyz} \equiv \frac{N_y|\mathcal{Y}|}{N_x|\mathcal{X}|\overline{m}_x\overline{m}_y}$ and $\overline{\theta}_{xyz} \equiv \frac{N_y|\mathcal{Y}|\overline{m}_x\overline{m}_y}{N_x|\mathcal{X}|}$.

Defining $\overline{\Theta} \equiv \bigtimes_{x,y,z}\left[\underline{\theta}_{xyz}, \overline{\theta}_{xyz}\right]$, we have that $\mathbf{T}\left(\overline{\Theta}\right) \subseteq \overline{\Theta}$. Since $\mathbf{T}$ is continuous on $\overline{\Theta}$, Brouwer's fixed point theorem implies that there exists at least one $\boldsymbol{\theta}^* \in \overline{\Theta}$ such that $\boldsymbol{\theta}^* = \mathbf{T}\left(\boldsymbol{\theta}^*\right)$.[33]

A.4 Proof of Proposition 2

We show uniqueness by leveraging Theorem 1 in Berry et al. (2013), henceforth denoted as BGH. Rewrite the system of equations in Equation (9) in logs, i.e.,

$$\log \theta_{xyz} = \log S^f_{xz|y}(\boldsymbol{\theta}) - \log S^w_{yz|x}(\boldsymbol{\theta}),$$

and define the mapping $\boldsymbol{\sigma} : \Theta \to \mathbb{R}^{|\mathcal{X}|\times|\mathcal{Y}|\times|\mathcal{Z}|}$ with typical element

$$\sigma_{xyz}(\boldsymbol{\theta}) \equiv \log \theta_{xyz} + \log S^w_{yz|x}(\boldsymbol{\theta}) - \log S^f_{xz|y}(\boldsymbol{\theta}). \tag{A.3}$$

Directed search equilibria $\boldsymbol{\theta}$ coincide with the roots of $\boldsymbol{\sigma}$. We further define the auxiliary mapping, with 0 a fictitious outside submarket,

$$\sigma_0(\boldsymbol{\theta}) \equiv 1 - \sum_{x,y,z} \sigma_{xyz}(\boldsymbol{\theta}). \tag{A.4}$$

We first establish two auxiliary lemmas about these mappings.

Lemma A.3. *For all (x, y, z) and $(x', y', z') \neq (x, y, z)$, and for all $\boldsymbol{\theta} \in \Theta$,*

$$\frac{\partial \sigma_{xyz}(\boldsymbol{\theta})}{\partial \theta_{x'y'z'}} \begin{cases} < 0, & \textit{if } x = x' \textit{ or } y = y', \\ = 0, & \textit{otherwise.} \end{cases}$$

Proof. By Equation (A.3) and Lemma A.2,

$$\frac{\partial \sigma_{xyz}(\boldsymbol{\theta})}{\partial \theta_{x'y'z'}} = \frac{\partial \log S^w_{yz|x}(\boldsymbol{\theta})}{\partial \theta_{x'y'z'}} - \frac{\partial \log S^f_{xz|y}(\boldsymbol{\theta})}{\partial \theta_{x'y'z'}}$$
$$\begin{cases} < 0, & \text{if } x = x' \text{ or } y = y', \\ = 0, & \text{otherwise.} \end{cases}$$

□

[33] The map $\mathbf{T}$ is continuous on $\bar{\Theta}$: $S^w_{yz|x}(\boldsymbol{\theta})$ and $S^f_{xz|y}(\boldsymbol{\theta})$ are continuous in $\boldsymbol{\theta}$, and $S^w_{yz|x}(\boldsymbol{\theta}) > 0$ throughout $\bar{\Theta}$ since the logit shares are strictly positive and $N_x > 0$.

Lemma A.4. *If a submarket* (x,y,z) *satisfies*

$$\pi^w_{yz|x}(\boldsymbol{\theta}) \leqslant \frac{1}{|\mathcal{Y}||\mathcal{Z}|}, \quad \textit{and} \quad \pi^f_{xz|y}(\boldsymbol{\theta}) \leqslant \frac{1}{|\mathcal{X}||\mathcal{Z}|},$$

then $\frac{\partial \sigma_0(\boldsymbol{\theta})}{\partial \theta_{xyz}} < 0$.

Proof. By Equations (A.3) and (A.4) and Lemma A.2,

$$\begin{aligned}
\frac{\partial \sigma_0(\boldsymbol{\theta})}{\partial \theta_{xyz}} &= -\sum_{x',y',z'} \frac{\partial \sigma_{x'y'z'}(\boldsymbol{\theta})}{\partial \theta_{xyz}} \\
&= -\frac{1}{\theta_{xyz}} - \sum_{y',z'} \frac{\partial \log S^w_{y'z'|x}(\boldsymbol{\theta})}{\partial \theta_{xyz}} + \sum_{x',z'} \frac{\partial \log S^f_{x'z'|y}(\boldsymbol{\theta})}{\partial \theta_{xyz}} \\
&= -\underbrace{\frac{1}{\theta_{xyz}}}_{>0} - \underbrace{\frac{\partial \overline{U}_{xyz}(\theta_{xyz})}{\partial \theta_{xyz}}}_{\geqslant 0} \underbrace{\left(1 - |\mathcal{Y}||\mathcal{Z}|\pi^w_{yz|x}(\boldsymbol{\theta})\right)}_{\geqslant 0} + \underbrace{\frac{\partial \overline{\Pi}_{xyz}(\theta_{xyz})}{\partial \theta_{xyz}}}_{\leqslant 0} \underbrace{\left(1 - |\mathcal{X}||\mathcal{Z}|\pi^f_{xz|y}(\boldsymbol{\theta})\right)}_{\geqslant 0} \\
&< 0.
\end{aligned}$$

□

We now verify that Assumptions 1–3 of BGH are satisfied. Theorem 1 of BGH then implies that $\boldsymbol{\sigma}$ is inverse isotone, hence injective, so the equilibrium system has at most one root.

- Assumption 1 of BGH (domain)
  The domain $\overline{\Theta}$ is a Cartesian product (see the proof of Proposition 1).
- Assumption 2 of BGH (weak substitutes)
  By Lemma A.3, $\frac{\partial \sigma_{xyz}(\boldsymbol{\theta})}{\partial \theta_{x'y'z'}} \leqslant 0$ for all (x,y,z) and $(x',y',z') \neq (x,y,z)$.
- Assumption 3 of BGH (connected strict substitution)
  By Lemmas A.3 and A.4, there is a directed path of strict substitution from each submarket to the auxiliary submarket 0. First, Lemma A.3 ensures a directed path from any submarket (x,y,z) to any other (x',y',z'):

  $$(x,y,z) \longrightarrow (x,y',z') \longrightarrow (x',y',z'),$$

  sharing the worker type x on the first edge and the firm type y' on the second. Second, since (x',y',z') was arbitrary, take it to be a submarket for which the conditions of Lemma A.4 hold. For this submarket there is a directed path

  $$(x',y',z') \longrightarrow 0.$$

  Together these yield a directed path of strict substitution from every submarket to the auxiliary submarket.

Finally, we provide economically interpretable conditions under which the inequalities in Lemma A.4 hold. The condition requires a single submarket (x, y, z) that is unattractive enough to both sides, relative to the other options available to a worker of type x or a firm of type y, that neither side would direct search to it more often than under the uniform-search benchmark.

Lemma A.5. *For all* (x, y, z) *and* $\boldsymbol{\theta} \in \Theta$,

$$\overline{U}_{xyz}(\theta_{xyz}) \leqslant \frac{1}{|\mathcal{Y}||\mathcal{Z}|} \sum_{y',z'} \overline{U}_{xy'z'}(\theta_{xy'z'}) \implies \pi^{w}_{yz|x}(\boldsymbol{\theta}) \leqslant \frac{1}{|\mathcal{Y}||\mathcal{Z}|},$$

$$\overline{\Pi}_{xyz}(\theta_{xyz}) \leqslant \frac{1}{|\mathcal{X}||\mathcal{Z}|} \sum_{x',z'} \overline{\Pi}_{x'yz'}(\theta_{x'yz'}) \implies \pi^{f}_{xz|y}(\boldsymbol{\theta}) \leqslant \frac{1}{|\mathcal{X}||\mathcal{Z}|}.$$

Proof. We prove the worker-side implication; the firm-side argument is symmetric. Applying the exponential function to both sides of the inequality, we obtain

$$\exp\left(\overline{U}_{xyz}(\theta_{xyz})\right) \leqslant \exp\left(\frac{1}{|\mathcal{Y}||\mathcal{Z}|} \sum_{y',z'} \overline{U}_{xy'z'}(\theta_{xy'z'})\right).$$

Since the exponential function is convex, by Jensen's inequality,

$$\exp\left(\frac{1}{|\mathcal{Y}||\mathcal{Z}|} \sum_{y',z'} \overline{U}_{xy'z'}(\theta_{xy'z'})\right) \leqslant \frac{1}{|\mathcal{Y}||\mathcal{Z}|} \sum_{y',z'} \exp\left(\overline{U}_{xy'z'}(\theta_{xy'z'})\right).$$

Together, this gives

$$\exp\left(\overline{U}_{xyz}(\theta_{xyz})\right) \leqslant \frac{1}{|\mathcal{Y}||\mathcal{Z}|} \sum_{y',z'} \exp\left(\overline{U}_{xy'z'}(\theta_{xy'z'})\right).$$

Dividing both sides by $\sum_{y',z'} \exp\left(\overline{U}_{xy'z'}(\theta_{xy'z'})\right) > 0$ and using the definition of $\pi^{w}_{yz|x}(\boldsymbol{\theta})$ yields the conclusion. □

A.5 Derivation of Equation (16)

Using Equation (5), we have that

$$\Delta_z W_y = \sum_x \mu_x \left[\left((1 - F(\theta_{xyz}))\rho_{xyz} - F(\theta_{xyz})\alpha_{xyz}\right) - \left((1 - F(\theta_{xyz'}))\rho_{xyz'} - F(\theta_{xyz'})\alpha_{xyz'}\right)\right].$$

Adding and subtracting both $\sum_x \mu_x((1 - F(\theta_{xyz}))\rho_{xyz'}$ and $\sum_x \mu_x F(\theta_{xyz}))\alpha_{xyz'}$ yields Equation (16).

A.6 Derivation of Equation (17)

We have that

$$\begin{aligned}
\Delta_x G &\equiv \sum_{y,z} \left(M_{xyz} W_{xyz} - M_{x'yz} W_{x'yz}\right) \\
&= \sum_{y,z} \left((W_{xyz} - W_{x'yz}) M_{xyz} + (M_{xyz} - M_{x'yz}) W_{x'yz}\right),
\end{aligned}$$

where the last equality follows from adding and subtracting $\sum_{y,z} M_{xyz} W_{x'yz}$. Using Equation (5), we obtain

$$\begin{aligned} W_{xyz} - W_{x'yz} &= (1 - F(\theta_{xyz}))\rho_{xyz} - F(\theta_{xyz})\alpha_{xyz}) - ((1 - F(\theta_{x'yz}))\rho_{x'yz} - F(\theta_{x'yz})\alpha_{x'yz}) \\ &= (\rho_{xyz} - \rho_{x'yz})(1 - F(\theta_{xyz})) + (\alpha_{x'yz} - \alpha_{xyz})F(\theta_{xyz}) + (F(\theta_{x'yz}) - F(\theta_{xyz}))\Phi_{x'yz}, \end{aligned}$$

where the last equality follows from adding and subtracting both $(1 - F(\theta_{xyz}))\rho_{x'yz}$ and $(1 - F(\theta_{xyz}))\alpha_{x'yz}$. In addition, one can verify that

$$M_{xyz} - M_{x'yz} = (M_{xy} - M_{x'y})\sigma_{x'yz} + (\sigma_{xyz} - \sigma_{x'yz})M_{xy}.$$

Together, this yields Equation (17).

B Application

B.1 Data

B.1.1 Occupation Types

This section provides an overview of the individual occupations, as defined by their three-digit SOC codes, that are included within each occupation type. Note that the individual occupations are the same as those selected in Goldin (2014).

Business.

- *Business Ops*
  Buyers and Purchasing Agents, Claims Adjusters, Appraisers, Examiners, and Investigators, Compensation, Benefits, and Job Analysis Specialists, Compliance Officers, Cost Estimators, Human Resources Workers, Logisticians and Project Management Specialists, Management Analysts, Market Research Analysts and Marketing Specialists, Meeting, Convention, and Event Planners, Miscellaneous Business Operations Specialists, Public Relations Specialists

- *Exec's and GMs*
  Chief Executives, General and Operations Managers

- *Financial Ops*
  Accountants and Auditors, Actuaries, Budget Analysts, Credit Analysts, Credit Counselors and Loan Officers, Financial Analysts and Advisors, Financial Examiners, Miscellaneous Financial Specialists

- *Lawyers and Judges*
  Judges, Magistrates, and Other Judicial Workers, Lawyers and Judicial Law Clerks

- *Sales*
  Advertising Sales Agents, First-Line Supervisors Non-Retail Sales, Insurance Sales Agents, Miscellaneous Sales Representatives, Services, Sales Engineers, Sales Representatives, Wholesale and Manufacturing, Sales and Related Workers, All Other, Securities, Commodities, and Financial Services Sales Agents

- *Managers*
  Administrative Services and Facilities Managers, Advertising and Promotions Managers, Architectural and Engineering Managers, Compensation and Benefits Managers, Computer and Information Systems Managers, Construction Managers, Economists, Emergency Management Directors, Financial Managers, Human Resources Managers, Industrial Production Managers, Marketing and Sales

Managers, Medical and Health Services Managers, Miscellaneous Managers, Public Relations and Fundraising Managers, Purchasing Managers, Training and Development Managers, Transportation, Storage, and Distribution Managers

Health.

- *Adv. Clin. Prof.*
  Nurse Anesthetists, Nurse Practitioners, Physician Assistants, Veterinarians
- *Clin. Prof.*
  Audiologists, Dietitians and Nutritionists, Miscellaneous Healthcare Diagnosing or Treating Practitioners, Psychologists, Registered Nurses, Therapists
- *Medical Practitioners*
  Chiropractors, Dentists, Optometrists, Pharmacists, Physicians, Podiatrists, Surgeons

Other.

- *Training & Post-Secondary Teachers*
  Education and Childcare Administrators, Post-Secondary Teachers, Training and Development Specialists
- *Other*
  Actors, Producers, and Directors, Agents and Business Managers of Artists, Performers, and Athletes, Air Traffic Controllers and Airfield Operations Specialists, Aircraft Pilots and Flight Engineers, Detectives and Criminal Investigators, First-Line Supervisors of Firefighting and Prevention Workers, First-Line Supervisors of Law Enforcement Workers, Ship and Boat Captains and Operators, Writers and Editors

Science.

- *Natural Scientists*
  Agricultural and Food Scientists, Astronomers and Physicists, Atmospheric and Space Scientists, Biological Scientists, Chemists and Materials Scientists, Environmental Scientists and Geoscientists, Miscellaneous Life Scientists, Miscellaneous Mathematical Science Occupations, Miscellaneous Physical Scientists, Natural Sciences Managers, Operations Research Analysts
- *Social Scientists*
  Architects, Except Naval, Miscellaneous Social Scientists and Related Workers, Surveyors, Cartographers, and Photogrammetrists, Urban and Regional Planners

Tech.

- *Comp. Programmers*
  Computer and Information Analysts, Software Developers and Computer Programmers
- *Cybersec. & Infrastr.*
  Database and Network Administrators and Architects, Information Security Analysts, Miscellaneous Computer Occupations
- *Engineers*
  Aerospace Engineers, Bioengineers and Biomedical Engineers, Chemical Engineers, Civil Engineers, Computer Hardware Engineers, Electrical and Electronics Engineers, Environmental Engineers, Industrial Engineers, Including Health and Safety, Marine Engineers and Naval Architects, Materials Engineers, Mechanical Engineers, Mining, Petroleum, and Geological Engineers, Miscellaneous Engineers

B.1.2 Data Sources

ACS. The American Community Survey (ACS) is an ongoing annual survey conducted by the United States Census Bureau. It collects detailed information on social, economic, demographic characteristics that is representative for the US population. The Census Bureau samples about 1% of the population each year, which yields approximately 3.5 million individuals. The Integrated Public Use Microdata Series (IPUMS) provides harmonized microdata versions of these data, ensuring that variables are coded consistently across survey years (Ruggles et al., 2025).

NLx Research Hub. The National Labor Exchange (NLx) is a non-profit online labor exchange operated by the National Association of State Workforce Agencies (NASWA), which represents the workforce agencies of all states in the United States, and by the DirectEmployers Association (DE). The exchange aggregates job vacancies from various individual sources, such as companies' career websites, states' job vacancy banks, and the federal jobs portal. These vacancies identify several characteristics of the job. We use information on the timing, region of employment, and occupation to construct the regional distributions of occupational vacancies (see below).

The vacancies are vetted to omit junk records and remove duplicates. The Research Hub of the NLx provides data, both real-time and historical snapshots, about all the job vacancies available on the exchange. The historical snapshots offer information on all open vacancies on any given day, and the NLx provides pooled monthly versions.

The NLx aims to cover all online job postings in the United States. But not all job openings are posted as online vacancies. Previous work by Hershbein and Kahn (2018), Deming and Noray (2020), and Acemoglu et al. (2022) has assessed how representative online vacancies are, using a commercial dataset by lightcast.io. They found that online vacancies are generally representative in terms of industries and occupations. The main shortcoming, see, for example, Section 3.A in Hershbein and Kahn (2018), is that highly skilled occupations tend to be overrepresented. Since our application focuses precisely on this segment of the labor market, this limitation is unlikely to pose serious concerns.

JOLTS. The Job Openings and Labor Turnover Survey (JOLTS) is collected by the Bureau of Labor Statistics (BLS) in the United States. The survey questions more than 20,000 establishments about their job openings and hires each month. Based on these questions, the BLS estimates the total number of job openings in the United States. For a vacancy to be measured in JOLTS, three conditions need to be satisfied on the last business day of the month: (i) a position exists and has work available, (ii) the job can start within thirty days, and (iii) there is active recruiting for workers from outside of the establishment.

Vacancy coverage. A comparison between the vacancies in JOLTS and NLx is provided at nlxresearchhub.org/nlx-joe. They highlight the substantial coverage of NLx vacancies. About 46.3% of JOLTS vacancies are represented in an NLx opening. And this is the aggregate number across all sectors and occupations. Coverage in our subset of highly skilled occupations is even higher. The two series also move similarly over time, with a correlation between the monthly openings in both datasets of 0.94.

Vacancy construction. To construct and scale the distribution of job vacancies across occupations for each state and year, we draw on two data sources: the first provides the total number of vacancies, and the second the occupational distribution.

- *Aggregate levels*
  The JOLTS also reports the number of job openings *per unemployed worker*. This allows for simple scaling with our sample.

- *Occupational distribution*
  To construct the job vacancy distributions, we keep all unique vacancies in NLx that open in a given month. We then pool all months and years to estimate the *distribution* of job vacancies at the state $\times$ occupation level. Vacancies which are open for several months are counted multiple times, as in JOLTS. While we could filter these out, including them likely yields a more accurate representation of

the vacancies that workers face a given moment within the year. We scale these distributions by the number of job vacancies per unemployed worker from JOLTS and the number of unemployed workers in our ACS sample.

B.2 Empirical Implementation

Normalization of parameters. As discussed in Section 3, our parametrization of α_{xyz} and ρ_{xyz} requires additional normalizations. Consider first the amenity parameters. We normalize the hours-related components to zero at long hours, setting $\alpha_{xz} = 0$ for every worker type x and $\alpha_{yz} = 0$ for every occupation y when $z =$ long hours. A further indeterminacy remains, as the levels of α_{xz} and α_{yz} at regular hours are identified only up to a common shift. We resolve it by additionally setting $\alpha_{yz} = 0$ for a single occupation at regular hours.[34] Together, these restrictions imply that the estimated amenity parameters measure the relative utility benefit of working regular over long hours. Consider next the productivity parameters. Symmetrically, we set $\rho_{yz} = 0$ for all occupations y at regular hours, so that the estimated productivity parameters measure the relative gain of working long over regular hours. For parsimony, we further set $\rho_{xz} = 0$ for all x and z.

Location and scale of earnings. Our identification treats the four macro-regions of the U.S. as cross-sectional draws from a common labor market that share the same surplus structure and matching technology. This assumption allows compositional variation in N_x and N_y across regions to identify the model. In practice, however, there are substantial regional wage differentials, reflecting differences in the cost of living, regional price levels, and broader local economic conditions that our model does not capture. To accommodate these, we include region fixed effects in the wage equation. In addition, because the model has no intrinsic unit for the surplus, we include a scale parameter for wages: concretely, we estimate an additional multiplicative constant in the wage function, leaving the units in which wages are measured unrestricted.

Scaling of matches. The data sources underlying our moments differ in magnitude: the matching patterns reflect the stock of existing employment relationships, whereas the unemployment and vacancy series are substantially smaller. To place these objects on a common basis, we scale all matching-pattern moments by the average monthly nonfarm separation rate over the sample period, as reported in JOLTS. This brings the matching moments to a magnitude comparable to the unemployment and vacancy

[34] We use *Social Scientists* for this normalization, since a preliminary fit of the model yielded the lowest surplus for this occupation. Setting $\alpha_{yz} = 0$ here is therefore unlikely to place the normalization in a region of the parameter space where surplus is most informative about the remaining parameters.

stocks, so that the implied volume of new matches is consistent with observed monthly hiring and separation flows.

B.3 Results

B.3.1 Descriptive Statistics

We first provide a brief overview of the sample, and show that that hours and earnings differ substantially across types. Next, we highlight differences in the occupational distributions of filled and vacant positions.

Table B.1 presents further summary statistics about the estimation sample, which is constructed from the 2022 to 2024 waves of the ACS. The sample is constructed similar to Goldin (2014), consisting of over 800,000 individuals individuals ages 25 to 64 that work full-time ($\geqslant$ 35 hours) and full year ($\geqslant$ 40 weeks). We have slightly more men ($\approx$ 53.3%) than women, which is because we focus on selected occupations. The unemployment rate is in line with the aggregate number ($\approx$ 4.1%). All individuals in our sample hold at least a BA degree, and 44% have a MA$^+$. Women are slightly higher educated than men, while men are significantly more likely to have a STEM undergraduate degree (43 vs 20%). The remaining 57 and 80% instead studied AHSS fields. About 58% are 40 years or older, with men slightly older than women.

The occupations of men and women are relatively similar. In total, about 54% work in *Business*, 15% in *Health*, 18% in *Tech*, 6% in *Science*, and 8% in *Other* occupations. The main difference is that women work more in *Health* (23% compared to 7%) and men work more in *Tech* (26% compared to 8%). We also find that men (26%) are more likely than women (19%) to work long hours. The raw gender gap in annual earnings is just over 20%, with men earning about 140,000 dollars per year compared to women's 110,000 dollars before taxes. There is also more variance in men's earnings than in women's. The gender gap in hourly wages is just below 20%, with men earning 60 dollars compared to women's 50 dollars, again with higher variance among men.

Hours and earnings. Figure B.1 shows the variation in hours and annual earnings across worker- and occupation types. Panel (a) highlights earnings differences across occupations. The highest earning groups in the sample are *Medical Practitioners*, *Executives and General Managers*, and *Lawyers and Judges*, who all earn between 150,000 and 200,000 dollars per year. Just below them are upper-level *Business* occupations, such as *Managers*, *Sales*, and *Financial Operations*, as well as *Computer Programmers*, *Engineers*, and *Advanced Clinical Professionals*, who earn between 1100,000 and 150,000 dollars per year. At the lower end, *Clinical Professionals* and *Training and Teaching* occupations have average annual salaries of just below 100,000 dollars.

Table B.1: Summary Statistics

	Men	Women	Total
Observations			
Employed	416,113	361,667	777,780
Unemployed	15,471	17,216	32,687
Demographics			
Master's Degree +	0.42	0.46	0.44
Age $\geqslant 40$	0.59	0.56	0.58
STEM Field	0.43	0.2	0.32
Occupations			
Business			
Business Ops	0.09	0.13	0.11
Exec's and GMs	0.06	0.04	0.05
Financial Ops	0.06	0.08	0.07
Lawyers and Judges	0.03	0.03	0.03
Managers	0.22	0.21	0.21
Sales	0.08	0.06	0.07
Health			
Adv. Clin. Prof.	0.01	0.03	0.02
Clin. Prof.	0.03	0.16	0.09
Medical Practitioners	0.03	0.04	0.04
Other			
Other	0.03	0.02	0.02
Training & Teaching	0.05	0.08	0.06
Science			
Natural Scientists	0.05	0.05	0.05
Social Scientists	0.01	0.01	0.01
Tech			
Comp. Programmers	0.1	0.03	0.07
Cybersec. & Infrastr.	0.06	0.03	0.04
Engineers	0.1	0.02	0.07
Hours			
Regular Hours	0.74	0.81	0.77
Long Hours	0.26	0.19	0.23
Earnings			
Avg. Annual Earnings (in k$)	140.48	110.93	12.67
Std. Dev. Ann. Earnings (in k$)	88.57	68.19	8.11
Avg. Hr. Wage (in k$)	0.06	0.05	0.06
Std. Dev. Hr. Wage (in k$)	0.04	0.03	0.03

Notes. This table shows summary statistics for the ACS (2022–2024). The selected sample consists of 25 to 64 year old individuals with at least a Bachelor's degree that are either unemployed or employed full-time and full-year in a higher earning occupation as in Goldin (2014).

Panel (b) highlights differences in hours across the same occupation types. Working long hours is particularly common in the highest earning occupations, being *Execu-*

Figure B.1: Earnings and Hours by Type

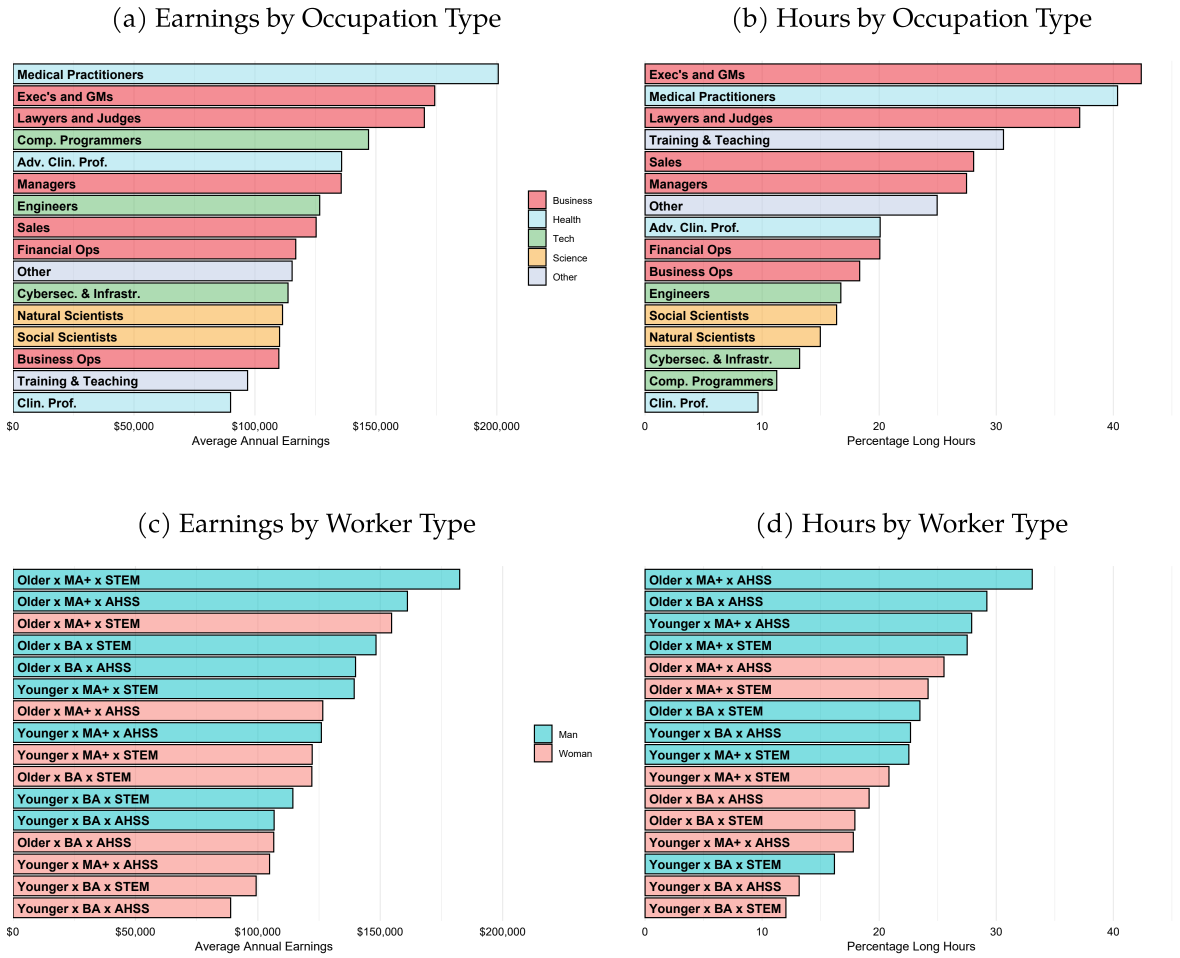

Notes. These figures present earnings and hours by worker and occupation type. Estimates are based on the ACS (2022–2024). Panels (a) and (c) show the distribution of annual earnings across the occupation and worker types defined in Section 4. Panels (b) and (d) show the fraction of workers that work long hours, defined as more than 50 per week, across the same types.

tives and General Managers, *Lawyers and Judges*, and *Medical Practitioners*. Among them, 35 to 45% reports working more than 50 hours per week. In lower earnings occupations, such as *Clinical Professionals*, *Social*, and *Natural Scientists*, working long hours is relatively unusual. Among them, around 15% or less reports long weekly hours. While higher earning occupations are generally characterized by longer hours, there are important exceptions. In *Tech* occupations only a small fraction of individuals reports long hours, yet their monthly earnings are among the highest, while *Training and Teaching* occupations show the opposite pattern.

Panel (c) shows earnings differences across worker types. This highlights that older men are the highest paid group in the sample, irrespective of education levels. Their average annual earnings range from about 140,000 (BA, AHSS) to over 175,000 dollars

(MA⁺, STEM). The only other group that earns more than 150,000 are older women with MA⁺ education and a STEM undergraduate degrees. The next groups are older women and younger men with MA⁺ education and an AHSS undergraduate field, generally earning between 110,000 and 125,000 dollars per year. The remaining groups, consisting mainly of younger men and women without graduate education, earn between about 80,000 and 110,000 dollars per year. These earnings differences by gender and education are well-documented, see for example the review by Blau and Kahn (2017).

Panel (d) highlights differences in working hours across worker types. The first main pattern is that, within education and experience levels, men are more likely to work long hours than women. This difference is particularly pronounced for older workers, among whom up to 30% of men with BA education work long hours, increasing to almost 35% among MA⁺ educated men who obtained their degree in an AHSS field. The corresponding percentages for women of the same type are closer to 20 and 25 respectively. A similar gender gap in hours among full-time workers is presented in Cortés and Pan (2019). There is also substantial heterogeneity by age and education. Younger women are very unlikely to work long hours, particularly if they obtained only a BA degree. Among them, only about 10% reports working long hours. Finally, we also see that individuals who graduated in a STEM field are generally less likely to work long hours than comparable workers that opted for an AHSS field.

Figure B.2: Hours within Worker and Occupation Types

(a) Within Occupation Type (b) Within Worker Type

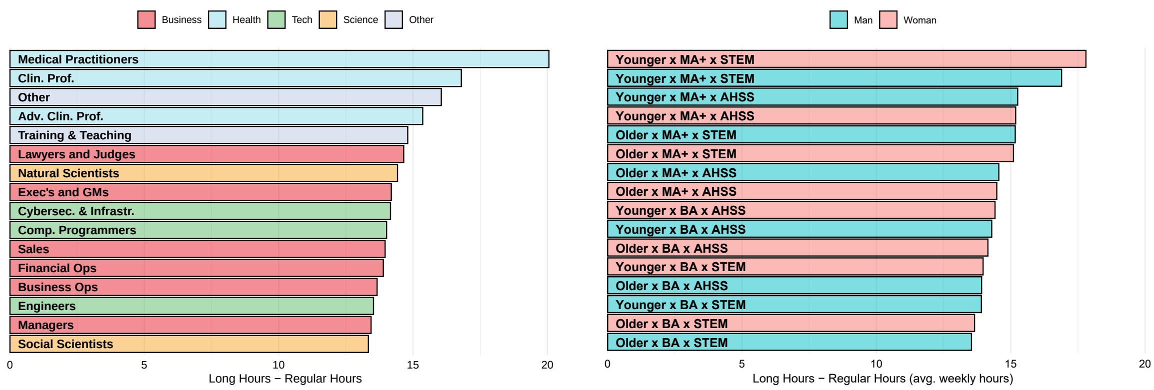

Notes. This figure highlights the average difference in the number of hours worked between long- and regular-hour positions. Panel (a) highlights the differences by occupation. Panel (b) shows the differences by worker type.

Hours within types. This Appendix highlights variation in working hours *within* worker and occupation types. Figure B.2 shows the difference in weekly hours between a long and regular hour position. To better understand this difference, note that average weekly hours in a regular position are about 40. Panel (b) highlights how regular- and long-hours jobs differ across occupation types. Across all occupations, the difference is close to 15 hours. The main exception is *Medical Practitioners*, who work about 20 hours per week more in long-hours positions. Panel (b) shows the same difference across worker types. There is again limited variation across types, and the average difference is close to 15 hours. The main exception here are younger younger men and women with STEM undergraduate degrees. For them, working long instead of regular-hours means an additional 17 hours per week.

Figure B.3: The Distributions of Filled and Vacant Jobs

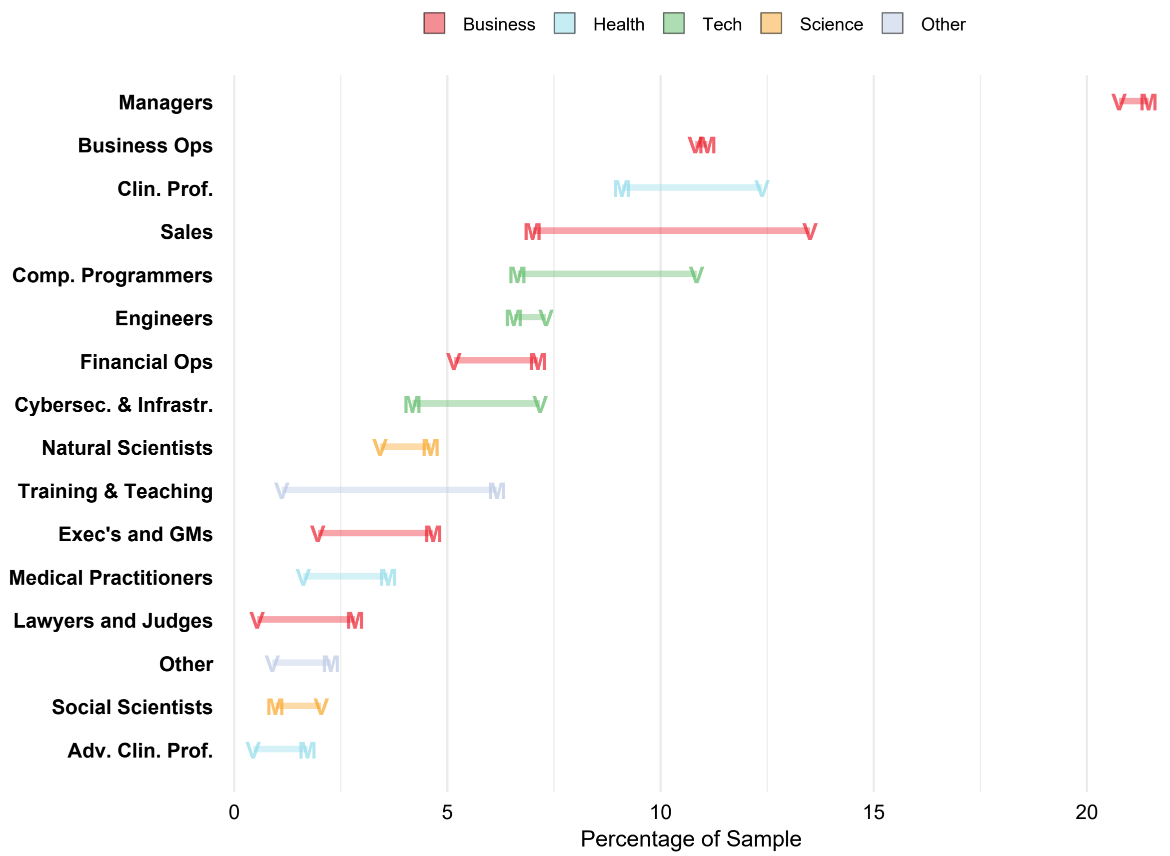

Notes. This figure highlights the occupational distributions of matched jobs (M) and vacancies (V). Constructed using ACS (matches) and NLx Research Hub (vacancies) datasets.

Filled and vacant positions. Figure B.3 examines the distributions of filled and vacant jobs across occupations. The first main pattern is the substantial variation in the relative size of different occupations. *Business* roles such as *Managers*, *Sales*, *Business-* and *Financial Operations*, account for a large share of both matches and vacancies. They jointly make up over 45% of both filled and vacant positions. In contrast, *Social Scientists*, *Advanced Clinical Professionals*, and *Other* occupations together represent less than 5% of observations in either group. The relative sizes of different occupations in *filled* positions aligns well with Cortés and Pan (2019). While there is little comparable evidence on *vacancies*, the distribution resembles that reported in the Appendix of Hershbein and Kahn (2018). Because they study a different period and use an alternative classification of occupations, the numbers are nonetheless difficult to compare. The vacancy distribution is also in line with BLS projections of occupational growth.[35]

A second notable pattern is the gap between the occupational distributions of vacant and filled jobs. Some occupations make up a substantially larger share of all vacancies than of filled positions, suggesting considerable excess demand. The main examples are *Clinical Professionals*, *Sales*, and all *Tech* occupations. While these together account for over half of all vacancies, they represent just a third of matches. In other occupations, we see the inverse pattern, as they make up a larger share of filled positions than of vacancies. These include high-earning occupations, such as *Financial Operations*, *Executives and General Managers*, *Lawyers and Judges*, *Medical Practitioners*, and lower earning *Training & Teaching* occupations, indicating a relative excess in supply. Together, these results highlight the importance of introducing information about the distribution of vacancies, which does not necessarily resemble that of filled jobs.

Variation in marginal distributions across markets. To identify the matching function elasticity, we rely on variation across markets, defined in the empirical application by the four U.S. macro-regions (Northeast, Midwest, South, and West). Figure B.4 highlights this variation. Panels (a) and (b) show differences in the marginal distributions of worker and occupation types, while panels (c) and (d) show differences in earnings; colors indicate deviations from the type mean. This variation is substantial, particularly in the marginal distributions of worker and occupation types. In several cases, markets with the largest share exceed those with the smallest share by more than 50%. For example, *Advanced Clinical Professionals* account for just 0.35% of observations in the South compared to 0.70% in the Northeast. Similarly, *Computer Programmers* make up 12.6% in the West but just 9.5% in the Midwest. The differences in wages are smaller but often still substantial. For example, *Medical Practitioners* earn about 156,000 dollars per year in the Northeast, compared to 184,000 in the Midwest.

B.3.2 Structural Estimates

This section presents the remaining structural parameter estimates omitted from the main text. Figure B.5 shows the productivity estimates by worker and occupation type,

[35]For example, the Bureau of Labor Statistics (2024) list of the fastest growing occupations includes *Nurses*, *Data Scientists*, and *Information Security Analysts*, for which we document a large number of vacancies.

Figure B.4: Workers and Earnings across Markets

Notes. This figure shows the marginal distributions of worker and occupation types (Panels (b) and (a)), and their average earnings (Panels (c) and (d)) across the four macro-regions. Data from the ACS (2022–2024), JOLTS, and the NLx Research Hub job vacancy data. The distributions of worker and occupation types are presented as a fraction of observations within the macro-region, while earnings are in levels (k$). Colors indicate relative deviations from the regional average.

as measured by ρ_{xy}. Figure B.6 shows the amenity estimates by worker and occupation type, as measured by α_{xy}. Figure B.7 shows the amenity estimates by occupation type and hours, as measured by α_{yz}.

B.3.3 Decompositions

Women tend to sort into slightly tighter submarkets, thereby increasing their relative share of the surplus. Figure B.8 explores this further. Panel (a) presents heterogeneity in market tightness by worker type. It reveals a greater "shortage" of women with STEM majors than of men with the same major. Panel (b) shows in which occupations the shortage of women relative to men is most pronounced, as captured by the ratio of tightness across the two markets. Shortages are largest in *Tech* occupations such as *Computer Programmers* and *Cybersecurity & Infrastructure*, both of which are occupations where wages are driven in large part by bargaining power (see Table 8).

B.3.4 Counterfactual Policy Experiments

Experiment 1: Closing the firm-side gender gap. Figure B.9 shows the re-sorting patterns by gender across occupations. Figure B.10 shows that tightness rises for women in almost all occupations and falls for men in nearly all of them.

Experiment 2: Closing the gender gap in STEM. Figure B.11 displays the effects on average earnings by occupation as a function of the relative increase in supply of women with a STEM degree after the compositional change. There is clearly a negative relationship, as larger increases in supply of STEM women push down average earnings. STEM women tend to resort themselves into occupations such as *Advanced Clinical Professionals*, *Natural Scientists* and *Medical Practitioners*, which tends to decrease the tightness and thereby average earnings in these occupations. In contrast, occupations towards the upper-left tend to either experience a net larger outflow of AHSS women than inflow of STEM women, which increases tightness and thereby bargaining power (e.g., *Training & Teaching*, see Figure B.12) or due to a compositional change in which STEM women increase the average productivity within those occupations (e.g., *Managers*).[36]

Figure B.13 highlights the spillover effect of an influx from STEM women (compared to the baseline) into occupations, which causes increased competition, and subsequently downward pressure on average earnings for AHSS men.

[36] This is clear from Figure B.5, where for *Managers*, and across other demographic characteristics, women with a STEM degree have higher estimated (perceived) productive values than their AHSS counterparts.

Figure B.5: Productivity Value by Worker and Occupation Type

(a) Women

Group	Occupation	Young × BA × AHSS	Old × BA × AHSS	Young × MA+ × AHSS	Old × MA+ × AHSS	Young × BA × STEM	Old × BA × STEM	Young × MA+ × STEM	Old × MA+ × STEM
Business	Exec's and GMs	59.4	66.1	63.7	70.4	61.3	68.1	53.6	72.7
Business	Lawyers and Judges	48.7	62.1	66.2	70.7	51.6	50.6	54.4	71.3
Business	Financial Ops	34.5	37.2	37.0	40.2	34.1	37.6	37.1	40.1
Business	Managers	26.1	28.9	27.8	31.4	26.5	30.2	28.7	33.6
Business	Business Ops	26.6	28.9	28.4	31.4	26.8	29.5	29.4	32.2
Business	Sales	15.1	17.4	15.9	18.9	13.7	16.6	14.4	17.9
Health	Adv. Clin. Prof.	52.0	63.1	64.1	67.4	47.7	52.9	64.2	67.6
Health	Medical Practitioners	30.5	35.3	38.0	46.4	28.5	33.5	39.7	50.8
Health	Clin. Prof.	18.8	20.8	19.1	21.6	16.7	18.1	17.6	19.2
Tech	Engineers	19.4	21.4	19.0	23.2	21.8	24.8	23.1	26.8
Tech	Comp. Programmers	14.2	16.4	16.9	19.6	18.2	19.9	20.4	21.6
Tech	Cybersec. & Infrastr.	14.0	16.8	15.9	17.8	15.3	18.0	16.8	19.3
Science	Natural Scientists	29.1	31.8	31.8	35.6	30.4	33.2	33.5	38.1
Science	Social Scientists	10.2	11.8	12.2	15.1	11.5	14.8	12.7	17.2
Other	Other	39.6	42.9	26.3	44.5	29.0	43.2	30.0	46.2
Other	Training & Teaching	18.0	31.4	20.4	35.9	16.1	19.8	19.0	36.4

(b) Men

Group	Occupation	Young × BA × AHSS	Old × BA × AHSS	Young × MA+ × AHSS	Old × MA+ × AHSS	Young × BA × STEM	Old × BA × STEM	Young × MA+ × STEM	Old × MA+ × STEM
Business	Exec's and GMs	61.0	70.0	67.5	74.2	63.3	71.8	51.5	77.0
Business	Lawyers and Judges	51.3	55.1	67.6	73.3	49.7	52.6	52.6	59.3
Business	Financial Ops	35.9	41.3	38.8	44.3	35.4	41.0	39.1	45.0
Business	Managers	26.8	31.6	29.8	34.3	28.3	33.5	31.4	37.0
Business	Business Ops	26.6	30.5	29.7	33.6	27.7	31.6	30.8	35.6
Business	Sales	16.7	21.2	18.3	21.9	15.8	20.1	17.0	21.7
Health	Adv. Clin. Prof.	55.0	56.2	65.0	68.8	51.5	52.1	63.2	69.9
Health	Medical Practitioners	30.9	33.4	39.7	50.8	27.3	30.6	41.7	55.5
Health	Clin. Prof.	18.5	20.1	18.1	21.1	16.5	19.0	17.8	15.6
Tech	Engineers	20.3	23.9	21.6	25.5	23.4	27.2	25.3	30.0
Tech	Comp. Programmers	16.1	20.4	19.8	22.5	21.4	24.4	23.7	26.2
Tech	Cybersec. & Infrastr.	15.3	19.6	17.3	21.7	17.4	20.9	19.1	23.3
Science	Natural Scientists	29.5	33.3	32.9	37.6	30.8	34.4	34.7	40.4
Science	Social Scientists	10.4	15.5	12.9	17.0	13.2	18.8	14.3	19.9
Other	Other	40.6	47.3	27.4	48.0	28.6	49.4	29.1	38.6
Other	Training & Teaching	18.6	21.1	18.7	36.9	16.3	21.9	17.7	38.8

Notes. This figure shows the estimated contributions to the productivity value across worker × occupation types (ρ_{xy}), independent from any occupation-wide or worker-wide contributions. All values are expressed in terms of annual earnings (in 10k$).

Figure B.6: Amenity Value by Worker and Occupation Type

(a) Women

Group	Occupation	Young × BA × AHSS	Old × BA × AHSS	Young × MA+ × AHSS	Old × MA+ × AHSS	Young × BA × STEM	Old × BA × STEM	Young × MA+ × STEM	Old × MA+ × STEM
Business	Exec's and GMs	14.6	14.6	17.0	18.5	16.3	16.4	7.1	17.2
Business	Financial Ops	8.3	10.9	10.9	13.1	9.6	11.4	9.4	12.4
Business	Business Ops	6.5	8.0	8.7	10.3	7.5	9.0	7.5	9.4
Business	Lawyers and Judges	2.8	12.3	16.2	18.1	2.4	3.6	3.7	16.9
Business	Managers	5.6	7.1	7.8	9.4	6.1	7.2	6.8	7.3
Business	Sales	0.9	2.0	0.8	1.8	0.8	1.1	0.5	0.5
Health	Adv. Clin. Prof.	5.9	12.8	19.4	22.2	5.3	2.6	21.2	23.4
Health	Medical Practitioners	4.4	5.5	9.5	5.5	2.7	10.2	11.1	4.4
Health	Clin. Prof.	3.4	5.2	6.8	8.2	4.1	5.5	6.1	7.6
Tech	Engineers	0.4	1.5	4.5	3.4	5.0	4.7	6.2	5.8
Tech	Cybersec. & Infrastr.	0.1	0.8	0.9	2.9	0.9	2.2	0.9	2.4
Tech	Comp. Programmers	0.0	0.1	0.2	0.3	0.1	1.2	0.2	2.1
Science	Natural Scientists	7.8	9.1	11.5	11.8	11.7	12.8	13.6	13.3
Science	Social Scientists	0.3	0.3	2.2	2.4	2.0	0.9	3.2	1.4
Other	Other	11.8	13.0	1.3	17.9	2.1	14.7	4.2	16.8
Other	Training & Teaching	0.4	12.8	3.0	17.5	2.2	2.9	3.3	16.8

(b) Men

Group	Occupation	Young × BA × AHSS	Old × BA × AHSS	Young × MA+ × AHSS	Old × MA+ × AHSS	Young × BA × STEM	Old × BA × STEM	Young × MA+ × STEM	Old × MA+ × STEM
Business	Exec's and GMs	13.1	14.2	16.3	16.2	14.0	14.3	2.2	15.8
Business	Financial Ops	7.1	8.1	11.6	9.5	8.4	7.5	9.2	8.5
Business	Business Ops	5.5	7.5	8.4	8.0	6.4	7.4	7.1	7.4
Business	Lawyers and Judges	1.5	1.7	17.5	16.7	2.5	1.9	2.1	1.8
Business	Managers	4.6	6.2	7.3	7.2	5.2	5.6	5.9	6.2
Business	Sales	0.3	0.8	0.3	0.4	0.5	0.4	0.3	0.3
Health	Adv. Clin. Prof.	0.9	1.3	18.1	19.1	1.4	1.9	20.5	20.0
Health	Medical Practitioners	3.0	2.2	9.3	1.4	2.1	1.5	9.9	0.7
Health	Clin. Prof.	0.1	3.8	6.3	6.0	0.6	2.4	4.5	0.9
Tech	Engineers	2.4	3.8	5.7	4.8	6.2	6.3	7.4	6.7
Tech	Cybersec. & Infrastr.	0.3	0.9	1.8	0.9	1.6	2.4	1.6	2.1
Tech	Comp. Programmers	0.1	0.1	0.3	0.3	0.2	0.4	0.3	1.5
Science	Natural Scientists	7.3	9.6	12.2	10.6	10.9	11.6	13.6	12.3
Science	Social Scientists	0.2	0.3	2.0	2.0	1.4	0.7	3.4	1.8
Other	Other	11.3	12.5	0.9	16.0	1.4	10.3	1.5	1.5
Other	Training & Teaching	0.6	0.8	1.5	16.3	1.1	0.8	1.5	15.6

Notes. This figure shows the estimated contributions to the amenity value across worker × occupation types (α_{xy}), independent from any occupation-wide or worker-wide contributions. All values are expressed in terms of annual earnings (in 10k$).

Figure B.7: Amenity Value by Hours and Occupation Type

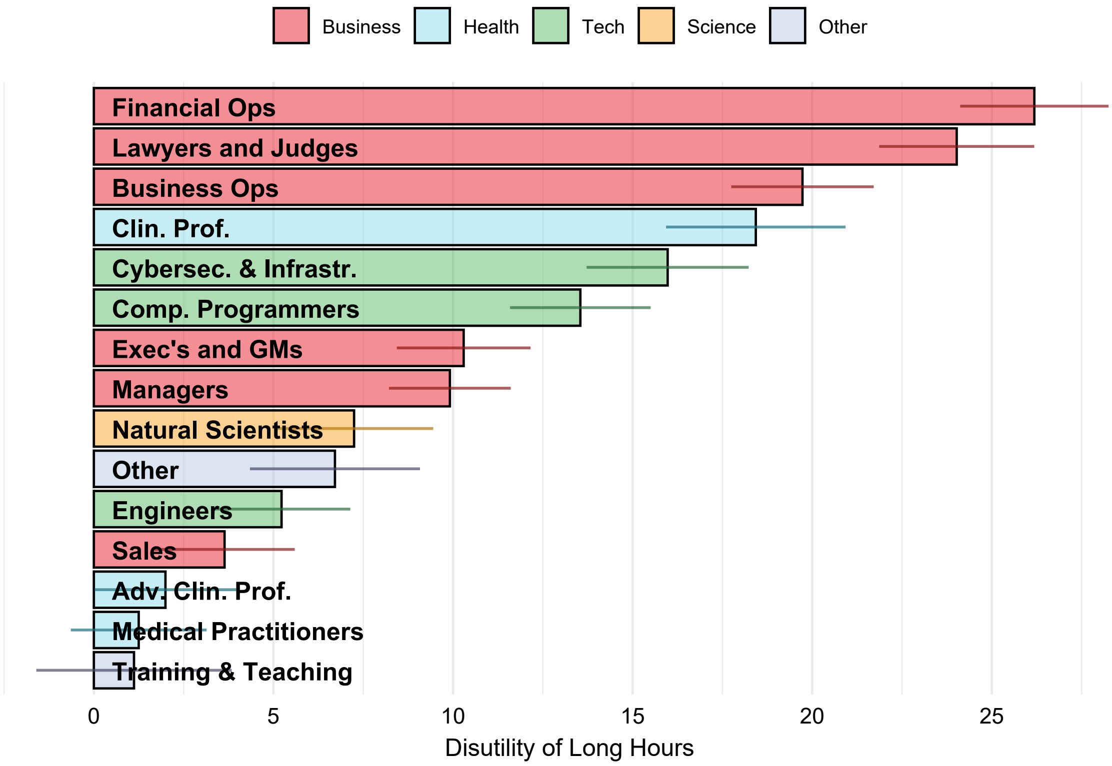

Notes. This figure shows the estimated utility premium for working regular instead of long weekly hours (α_{yz}), expressed as a percentage of mean earnings in long-hours jobs across occupations. Higher values indicate a greater utility from regular compared to long hours. The 95% confidence intervals are based on the asymptotic variance of the minimum distance estimator.

Figure B.8: Estimates: Market Tightness by Worker Type and Occupation

(a) By Worker Type

(b) By Occupation

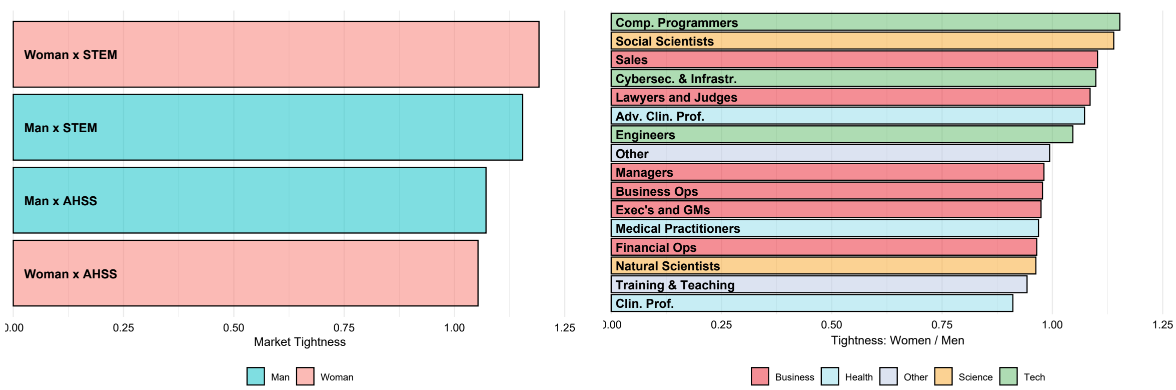

Notes. This figure shows the differences in tightness. Panel (a) highlights the heterogeneity of tightness by worker types, and Panel (b) heterogeneity by occupation.

Figure B.9: Occupational Resorting by Gender

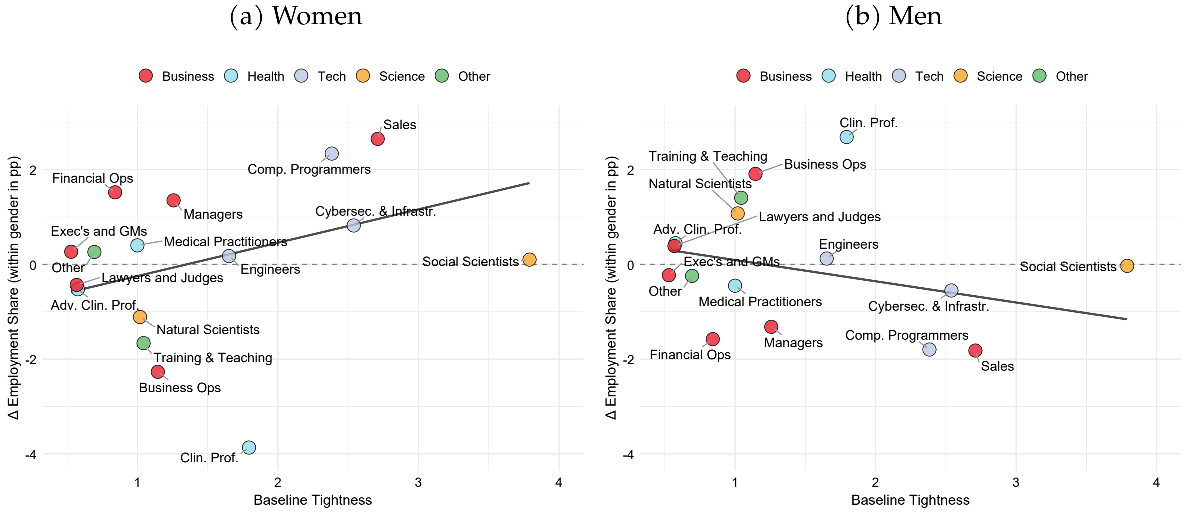

Notes. This figure plots the resorting patterns. Panel (a) shows resorting (in percentage point changes in employment shares) for women when the gender gap in firms' perceived productive value is fully closed ($\tau = 1$) on baseline tightness, and and Panel (b) shows resorting for men as a function of baseline tightness.

Figure B.10: Tightness Change by Gender and Occupation

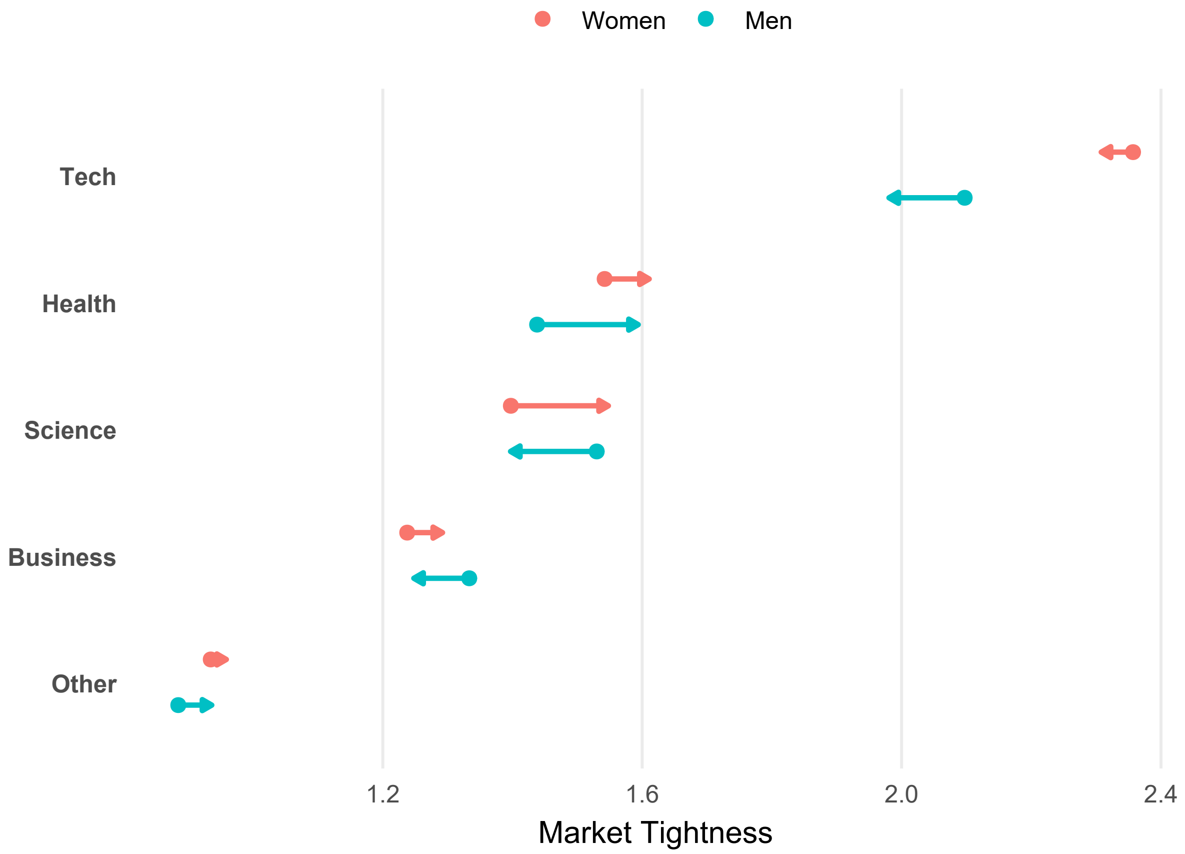

Notes. This figure plots the change in tightness from baseline to the counterfactual where the gender gap in firms' perceived productive value is fully closed ($\tau = 1$) across occupations.

Figure B.11: Estimates: Counterfactual Earnings and Sorting (Women)

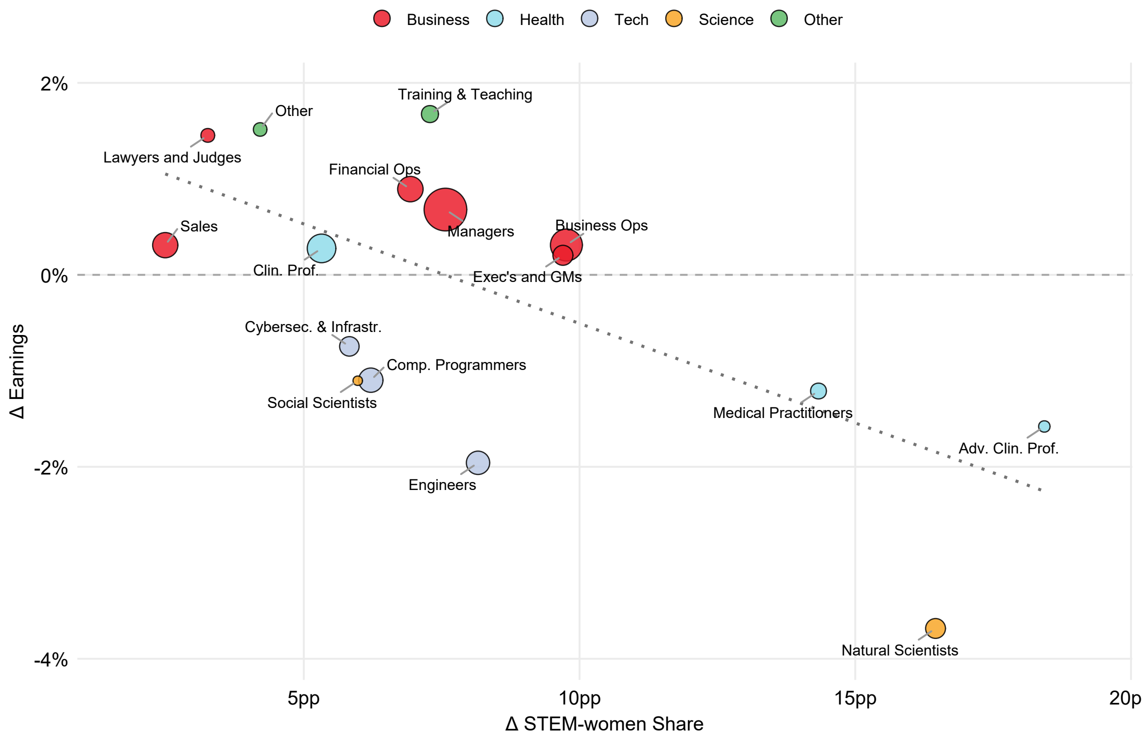

Notes. This figure highlights the relation between the difference in occupational earnings of women (y-axis) and the share of STEM women (x-axis) between the baseline and counterfactual equilibria. The bubble sizes indicate the total number of workers in each occupations in the baseline.

Figure B.12: Counterfactual Tightness by Occupation

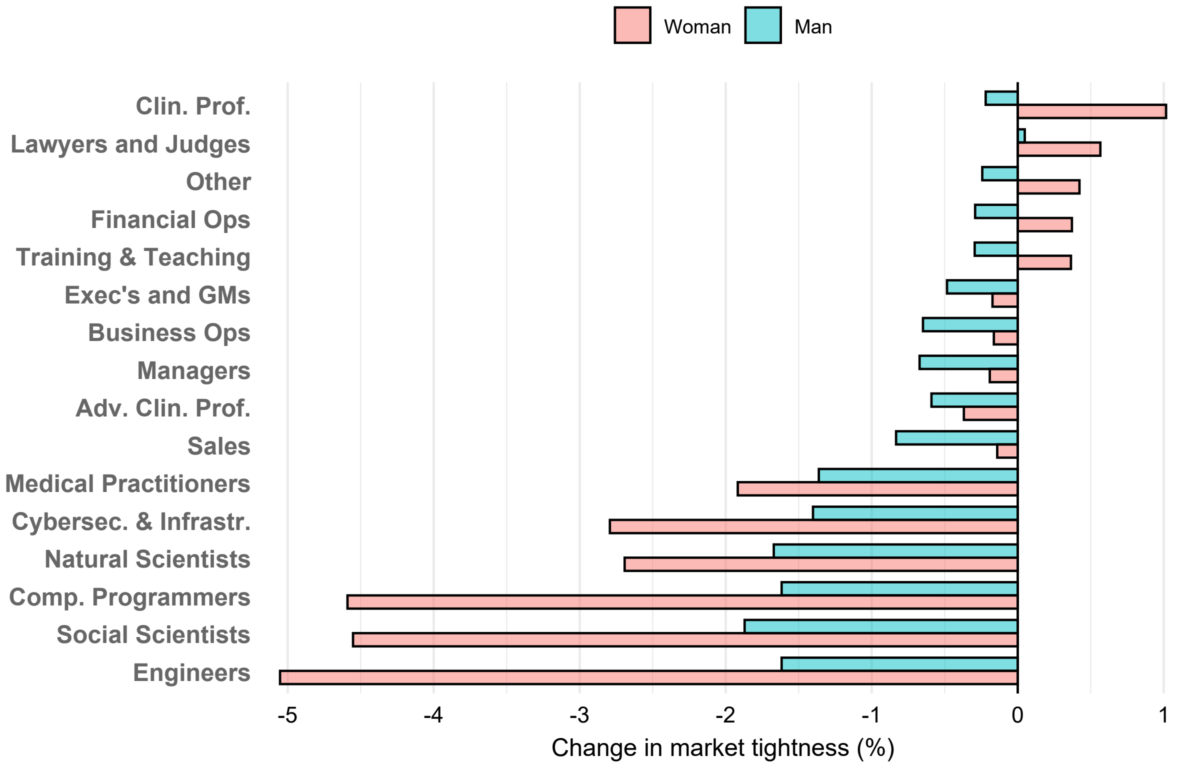

Notes. This figure plots the change in tightness from baseline to the counterfactual (across occupations, by gender) when we reallocate 20% of AHSS women to STEM.

Figure B.13: Counterfactual Earnings and Sorting (AHSS Men)

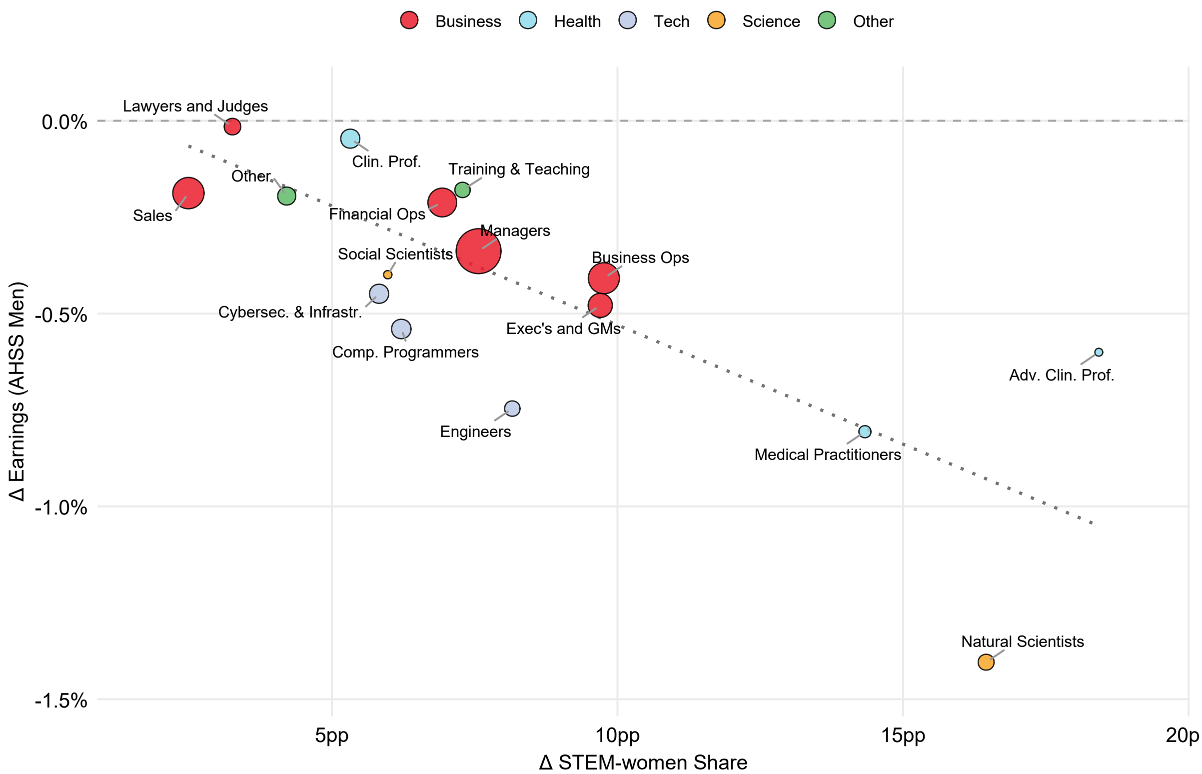

Notes. This figure highlights the relation between the difference in occupational earnings of AHSS men (y-axis) and the share of STEM women (x-axis) between the baseline and counterfactual equilibria. The bubble sizes indicate the total number of AHSS male workers in the baseline.