Workplace Amenities and Labor Market Inequality

Dana Scott*

January 21, 2026

Updated frequently. Click here for latest version.

Abstract

Jobs differ in wages and in the amenities they provide, particularly workplace flexibility—the temporal aspects of work including schedule predictability and demands on workers' time. While prior work documents wage-flexibility trade-offs, the role of production technology in shaping these compensating differentials remains largely unexplored. Using French matched employer-employee data linked to large-scale surveys, I document three empirical patterns: (1) flexibility co-varies within firms across establishments and occupations; (2) some flexibility dimensions correlate positively and others negatively with pay; and (3) women sort into jobs with better flexibility within occupations, suggesting preferences drive sorting. To explain these patterns, I develop and estimate an equilibrium model where firms choose wages and amenities subject to both worker preferences and technological constraints. The model reveals that high-skill occupations face productivity costs from providing flexibility roughly twice those of mid-skill occupations. Decomposing equilibrium outcomes, I find that technological parameters explain 22 percent of amenity variation. Counterfactual simulations demonstrate that equalizing women's amenity preferences to men's would reduce the gender wage gap by 4.6 percentage points (20 percent of the baseline gap), while mandating minimum flexibility levels would widen the gender gap by 1.8 percentage points. These findings reveal fundamental technological barriers to amenity provision that drive worker sorting and limit the effectiveness of flexibility mandates, with important implications for understanding the gender pay gap.

I am grateful to my advisors, Costas Meghir, Joe Altonji, and John Eric Humphries, for their continuous support and guidance throughout this project. I also thank Jason Abaluck, Steve Berry, Barbara Biasi, Richard Blundell, Pauline Carry, Guillermo Carranza Jordan, Crossan Cooper, Jingyi Cui, Janet Currie, Paul Goldsmith-Pinkham, Henrik Kleven, Bentley MacLeod, Victoria Marone, Derek Neal, Cormac O'Dea, Elin Sundberg, Winnie Van Dijk, and participants at the Yale Labor lunch for helpful discussions and comments. I gratefully acknowledge support from the Cowles Labor-Public Fund, the MacMillan Center International Dissertation Research Fellowship, and the Institute for Social and Policy Studies Graduate Fellowship. I thank the Institute for Fiscal Studies for generously hosting me as a visitor for access to the French administrative data.

*Department of Economics, Yale University. Email: dana.scott@yale.edu.

1. Introduction

Jobs differ in the wages they pay and in the amenities they provide. In particular, workplace flexibility—what Goldin (2014) defines as the temporal aspects of work, including when, where, and how predictably one works—plays a central role in workers' job choices and labor-market outcomes (Mas 2025; Mas and Pallais 2019; Maestas et al. 2023; Humlum et al. 2025). Workplace flexibility is especially relevant to gender inequality in the labor market. As past research has shown, women have a higher willingness to pay for this amenity than men (Goldin and Katz 2011; Goldin 2014; Mas and Pallais 2017; Wiswall and Zafar 2018). However, we know surprisingly little about *why* jobs provide different flexibility amenities. While the literature documents compensating differentials—wage-amenity trade-offs in equilibrium—it has largely overlooked how production technology shapes these trade-offs. Flexibility can directly affect worker productivity: for example, allowing workers to set their own schedules or sign off in the evening can hamper coordination. These technological costs constrain firms' flexibility choices, but the relative importance of production constraints versus worker preferences in determining equilibrium amenities remains an open question.

To what extent are amenities determined by firms' production technologies versus workers' preferences, and how does this impact labor market inequality? In this paper, I study how these two forces determine wages and amenities in labor market equilibrium, focusing on the role of workplace flexibility.[1] I study two specific dimensions of workplace flexibility: (1) scheduling irregularity, such as night shifts or variable hours, and (2) "greedy hours", or the ways in which jobs place demands on workers' time outside of the usual workday—such as teachers grading assignments in the evening after a full day at school. Using French administrative and survey data, I document systematic patterns suggesting that technological costs, not just worker sorting, drive compensating differentials. I then estimate an equilibrium model incorporating both worker preferences and firm-specific technological costs of providing flexibility. The model reveals substantial heterogeneity in these costs: high-skill occupations face productivity penalties roughly double those of mid-skill workers. Counterfactual analysis demonstrates that policies mandating minimum flexibility amenities face fundamental trade-offs: they improve welfare for workers in less competitive markets but widen the gender pay gap by deepening gendered sorting to jobs with different compensation bundles. These findings highlight how production technologies, not just preferences, shape the distribution of job quality and constrain the scope for policy interventions aimed at improving workplace flexibility and gender inequality.

Empirical work on the pay–amenity relationship has been limited by the difficulty of measuring

[1] Following Goldin (2014), I define workplace flexibility as temporal aspects of work: predictability of schedules, control over hours, and requirements to work evenings, nights, or be "on call." I use "flexibility" and "amenities" interchangeably throughout.

amenities. Unlike wages, which are systematically recorded in administrative data, non-pay amenities are rarely observed at scale, making them difficult to study. The French setting allows me to overcome these challenges by linking matched employer-employee data to two large-scale surveys to observe a rich set of characteristics related to job flexibility. The literature has typically addressed measurement challenges by inferring an aggregate "amenity value" from worker flows or treated amenities as residuals that explain firm size and hiring, without identifying what features make jobs attractive or why some firms provide good amenities and others do not (Sorkin 2018; Taber and Vejlin 2020; Lamadon et al. 2022; Lehmann 2025; Volpe 2025). However, this approach faces two key limitations. First, treating unobserved amenities as endogenous requires strong identifying assumptions. Second, without knowing what these amenities represent, we cannot assess their implications for inequality or policy. By observing flexibility amenities, I am able to address the measurement challenge, endogenize amenity choices in the firm's problem, and study how supply and demand forces contribute to amenity determination.

My analysis proceeds in three parts. In the first part, I develop three empirical facts on how firms set flexibility amenities. These facts suggest that technological and organizational costs, not just worker preferences, drive observed wage-amenity trade-offs.

The first fact shows that firms set amenities systematically: most of the variation in amenities occurs at the level of the firm, often across many establishments, with variation across occupations within the firm playing a secondary role. I regress a variety of observed amenities on different combinations and interactions of fixed effects for firms, occupations, and locations. The adjusted R-squared contribution of firm fixed effects alone to many of these models is on the order of 70-75 percent. For many amenities, interacting the firm and occupation fixed effects adds an important secondary source of explanatory power to the regression. If worker preferences alone drove amenity provision, firms hiring from multiple labor markets with distinct worker preferences would show substantial within-firm variation in amenities. Instead, the strong firm-level commonalities point to shared technological or organizational constraints as key drivers of amenity choices.

The second fact documents that better-paying jobs offer some good amenities and some disamenities, suggesting that technological requirements, not just preferences, shape what firms provide. I define high-paying jobs based on their wage distributions within their sector, following the methodology of Bonhomme et al. (2019). High-paying jobs offer some features that are better than those at low-paying firms: they offer more consistent schedules, less exposure to hazards, and cultures that make workers feel more respected. However, workers in these same jobs also report facing more stress and overwork and are more likely to work in the evenings. If substitution in compensation were the only motivation for firms to offer amenities, we would expect that all amenities would be worse in high-paying jobs. These contrasts indicate that wage–amenity trade-offs

alone cannot fully explain firms' amenity choices.

The third fact documents that jobs chosen by men and women tend to offer different amenities. The pattern is not explained by occupational choice, but rather by choices of jobs *within* occupations. Women tend to have jobs with better visibility and stability in scheduling and less night and evening work. These observed gender differences in scheduling amenities motivate a model that can disentangle workers' preferences for amenities from technological differences across firms to explain gendered sorting.

The facts raise three questions. Why do amenities co-vary within firms (Fact 1)? Why are some amenities positively correlated and others negatively correlated with pay (Fact 2)? And how much of gender sorting reflects worker preferences versus firm technology (Fact 3)? To organize thinking about these questions, I develop a stylized model of how firms choose wages and flexibility amenities. The model rationalizes the observed facts and clarifies which economic primitives govern equilibrium compensation packages, motivating the estimation approach and counterfactuals. In the stylized model, monopsonistically competitive firms use two differentiated occupations as imperfectly substitutable inputs into production. The firm chooses the wage and the level of a single flexibility amenity for each occupation to maximize profit. Providing flexibility can reduce worker productivity; the magnitude of the penalty may differ by occupation. The model generates closed-form equations for equilibrium wages and amenities in each job: wages are a job-specific markdown to the marginal product of the worker, and amenities depend on workers' relative preference for the amenity and the firm's cost to provide it.

Motivated by the descriptive facts and stylized model, the second part of the paper develops and estimates an equilibrium model of labor supply and labor demand. In the model, I consider two key dimensions of flexibility amenities: (i) demanding schedules, such as working night shifts or working a different shift every week, and (ii) demands on workers' time outside of ordinary hours, characterized by working in the evenings and working at home. The labor supply side of the model draws on the differentiated products literature in industrial organization to estimate labor supply elasticities with respect to wages and amenities. Workers face a discrete job choice problem with a nested logit structure conditional on assignment to a given labor market.[2] Within the labor market, the worker first chooses a 2-digit occupation. Then, conditional on occupation, she chooses a job. She derives utility from wages, amenities, and the commuting distance between her residence and workplace. I estimate the model in two steps using administrative matched employer-employee data from 2017-2023, excluding 2020, matched to the annual Labor Force Survey (EEC). In the first step,

[2]I consider the possibility that workers choose a labor market in an additional top level of the nested logit problem. However, I estimate a precise zero for the substitution elasticity across markets in this specification, which rejects this model structure in favor of one in which the market is assigned and the occupation and job are chosen conditional on the market. Moreover, 98 percent of job-to-job transitions in the Panel-BTS occur within the markets I define here, further suggesting that voluntary substitution across markets is extremely limited.

I sequentially recover substitution parameters and job-specific mean utilities for men and women at each level of the nesting structure. In the second step, I estimate the utility value of wages and amenities using Hausman-style instruments motivated by my descriptive facts on firmwide wage and amenity setting. I compute job-specific labor supply elasticities that depend on market shares, substitution parameters, and the utility value of wages and amenities.

The supply model yields precise estimates of labor supply elasticities with respect to both wages and amenities for men and women. Women have higher willingness to pay for scheduling amenities than men: women value a 1 percent increase in wages equally to a 0.36-0.59 percent increase in this amenity, versus men's 0.18-0.35 percent. However, women's job-specific labor supply elasticities are consistently *lower* than men's within the same markets and skill levels, despite the higher willingness to pay, because their choices are less responsive to changes in the amenity levels. This elasticity gap is approximately 40 percent for wages and 22 percent for the "greedy hours" (dis)amenity. The labor supply elasticities imply wage markdowns of 36-39 percent for men and 48-51 percent for women.

I next estimate the labor demand side by modeling firms' production functions with differentiated occupational inputs and an explicit technological role for flexibility amenities, conditional on the elasticities estimated on the labor supply side. The demand side generalizes the stylized model of the firm's wage and amenity choice to allow for many occupations, many amenities, and unobserved heterogeneity in worker skill. Firms have CES production functions over differentiated occupational inputs. The firm chooses wages and flexibility amenities to maximize profit, where flexibility directly affects worker productivity through occupation-specific penalties. The first-order conditions yield closed-form equations for optimal wage and amenity choices. I estimate the CES parameter governing substitutability of occupational inputs in the production function, and the cost to provide workplace flexibility to workers for four skill groups using two-step GMM.

The results reveal a sharp skill gradient in flexibility costs: the cost to provide flexibility in high-skill occupations is roughly twice that for intermediate professionals, with even smaller costs for lower-skill categories. Decomposing variation in the structural equations shows that technology parameters alone explain 22 percent of the amenity variation, while preferences for the amenity relative to wages explains close to none. This suggests that firm technology is a first-order driver of flexibility provision. Decomposing the sources of variation in the structural wage equation reveals a strong negative correlation between firm productivity and flexibility costs. Because high-productivity firms incur large flexibility costs, providing workplace flexibility at high-paying jobs strongly attenuates wage variance.

In the third part of the paper, I simulate counterfactual changes to model parameters to quantify how firm costs and worker preferences shape equilibrium amenities, wages, and gender inequality. Equalizing women's amenity preferences to men's would raise wages for both genders and reduce

the gender wage gap by 4.6 percentage points—about 20 percent of the baseline gap. Reducing flexibility costs for high-skill workers to match those for medium-skill workers (a cost reduction of 50-60 percent) would have a smaller effect on the gender pay gap but would substantially increase amenity provision, improving welfare by 1-3 percent across the distribution. In contrast, mandating a minimum flexibility level at the 20th percentile would improve welfare by 2 percentage points for the bottom half of the distribution but widen the gender pay gap by 1.8 percentage points. Together, these counterfactuals demonstrate how preferences and technological costs jointly determine equilibrium outcomes and reveal that policies mandating minimum amenities create trade-offs between welfare gains and increased gender inequality.

This paper contributes to two major strands of literature. First, I advance the literature on non-wage amenities and compensating differentials in the labor market, recently reviewed by Mas (2025). While the theoretical literature has long recognized that both worker preferences and firm production technologies determine equilibrium wage-amenity packages (Rosen 1986; Hwang et al. 1998), empirical work has struggled to separate these forces. Substantial evidence documents that amenities significantly influence workers' job choices (Rosen 1986; Hamermesh 1999; Pierce 2001; Hall and Mueller 2018; Mas and Pallais 2017, 2019; Sockin 2022; Maestas et al. 2023), and that amenities related to time use and job flexibility are especially important for women (Goldin 2014; Juhn and McCue 2017; Wiswall and Zafar 2018; Chen et al. 2019; Cubas et al. 2021, 2023; Le Barbanchon et al. 2021; Kim et al. 2024). Typical administrative data sets do not contain much information on non-wage characteristics. This measurement limitation has led much of the literature to define amenity value as a structural residual inferred from revealed preferences (Taber and Vejlin 2020; Sorkin 2018; Lamadon et al. 2022; Lachowska et al. 2023; Lamadon et al. 2024; Lehmann 2025). It is difficult to treat unobserved amenities as endogenous using these tools, requiring additional assumptions for identification. Similarly, unobserved amenities are difficult to interpret in the firm's problem. Because of these limitations, the endogenous decisions of firms to provide amenities remain understudied. Recent work has made progress by developing systematic amenity measures and linking them to administrative data (Sockin 2022; Lagos 2025; Caldwell et al. 2025; Humlum et al. 2025; Morchio and Moser 2025). These papers represent an important step forward: they collect some of the best available data on this topic and develop new ways to measure how workers value amenities. My paper builds on this work by explicitly modeling amenities as endogenous firm decisions, making it possible to identify the supply and demand forces driving amenity provision and study their implications for sorting and inequality in equilibrium.

Of these papers, only Morchio and Moser (2025) endogenize firms' amenity choices by modeling heterogeneous amenity costs in a search framework. My paper complements theirs in three ways. First, I microfound differences in firms' costs to provide amenities by estimating productivity

penalties to providing more favorable time use amenities across different firm and worker skill types. Second, I focus on time use amenities that they do not cover due to data limitations. Third, I allow for occupational heterogeneity and imperfect substitution between different types of workers, allowing me to learn about the relationship between firms' technologies and amenity provision costs.

Second, I contribute to the large literature on imperfect competition and monopsony in the labor market, recently reviewed by Azar and Marinescu (2024) and Kline (2025).[3] I build on a class of equilibrium models with monopsonistically competitive firms that adapt differentiated product demand models from the industrial organization literature (Card et al. 2018; Berger et al. 2022; Azar et al. 2022; Lamadon et al. 2022; Volpe 2025). I make three contributions. First, I apply the nonparametric labor market clustering approach of Nimczik (2023) and Jarosch et al. (2024) to estimate a nested logit model of labor supply—the first to combine these approaches.[4] While existing studies either assume workers consider markets defined by sector, occupation, or location, or rely on naturally restricted choice sets like job applications, I use revealed job-to-job transitions to infer which jobs workers actually consider as substitutes. This nonparametric clustering partitions thousands of jobs into tractable labor markets and provides direct evidence on rich substitution patterns across the entire economy. Second, I incorporate observed non-wage amenities and estimate their elasticities, adapting the endogenous product characteristics framework of Fan (2013) with instruments for both wages and amenities building on Azar et al. (2022). Third, I allow labor supply elasticities to vary by gender to examine preference heterogeneity while maintaining tractability.

This paper is organized as follows. Section 2 discusses the data sources, sample selection, and key definitions of amenities and labor markets. Section 3 presents the stylized model of firm wage and amenity setting and three descriptive facts. Section 4 develops and estimates the labor supply model. Section 5 generalizes the stylized model of firm conduct to estimate labor demand parameters. Section 6 presents results from counterfactual simulations. Section 7 concludes.

2. Data, Sample Selection, and Measurement of Amenities

2.1. Administrative Data

My main worker-level outcomes are derived from the French matched employer-employee records, the *Base Tous Salariés* (BTS), focusing on the period from 2017 to 2023 (excluding 2020). I use the *Postes* file of the BTS, where observations are at the job spell × year level. Firms report pre-

[3]See also Manning (2021), Sokolova and Sorensen (2021), and Card (2022) for further reviews.

[4]Fogel and Modenesi (2022) use a conceptually related but different model to classify jobs into nonparametric labor markets, built not on job-to-job transitions but on match probabilities inferred from coworker networks. Felix (2022) also uses job-to-job transitions to microfound a nested logit choice structure, but partitions jobs into markets by microregion × occupation cells rather than using nonparametric clusters.

and post-tax earnings, hours worked, spell dates, enterprise and establishment identifiers (SIREN and SIRET), contract type (full/part-time, fixed-term/open-ended), industry (NAF codes), 4-digit occupation (PCS-ESE), and worker age and gender. For exercises requiring panel data on workers across years and job spells, I use the 2013-2021 file of the Panel-BTS, a representative subsample that tracks workers over time using anonymized identifiers.

2.2. Measurement of Job Characteristics in Survey Data

2.2.1. Continuous Labor Force Survey (EEC)

The *Échantillon Emploi en Continu* (EEC) is the French administration of the annual EU Labor Force Survey, recording earnings, hours, occupations, and job characteristics for employed workers. The annual sample includes 75,000-110,000 individuals aged 25-55 over 2017-2023. The EEC contains establishment and enterprise codes (SIRET and SIREN), occupation (PCS-ESE), and location (département), allowing linkage to administrative data at the job level (enterprise × occupation × location × year). I aggregate individual survey responses to the job level using survey weights. Table 7 summarizes the amenity variables; Appendix Table D.1 provides variable names and question text.

TABLE 1. Job Characteristics Derived from Survey Data Sources

Amenity Category	**Key Variables**
Panel A: Labor Force Survey (EEC)	
Work schedule	Same weekly schedule, Variable schedule, Choose own hours
Work location	Work from home (>50% of time, some of the time)
Non-standard hours	Evening work, Night work
Panel B: Working Conditions Survey (CT-RPS—8 Principal Components)	
Schedule: variability	Schedule variation week to week
Schedule: visibility	advance knowledge of hours
Odd hours	Work in the evening or at night (sometimes, often)
Flexibility	Bring work home, can interrupt work when desired
Stress and overwork	Work/contact outside anticipated hours, bring work home
Workplace hazards	Lifting, standing, strenuous movements
Culture: respect	Receive respect, treated equitably
Culture: relationships	Tension in relationships, receive help from others

Notes: This table summarizes observed characteristics covered by survey questions in the Labor Force Survey (EEC, Panel A) and Survey of Working Conditions and Psychosocial Risks (CT-RPS, Panel B). Detailed descriptions of the survey questions corresponding to the measured variables are provided in Appendix D.2.

2.2.2. Recovery of latent amenities in the EEC with a factor model

This paper estimates labor supply elasticities related to job flexibility. No single survey question perfectly captures this continuous characteristic, so I develop a latent factor model to recover the underlying amenity dimensions from multiple noisy measures in the EEC, following Bai and Ng (2002). I include five variables: same schedule every week, work from home > 50% of the time, work from home some of the time, work evenings, and work nights. I find that a 3-factor model best balances fit and interpretability.

TABLE 2. Latent Amenity Factors in EEC Data

Variable	Demanding Schedule	Greedy Hours	Work From Home
Fixed schedule	-0.38	0.14	0.15
Work evenings	0.90	0.10	0.07
Work nights	0.60	-0.21	-0.16
> 50% WFH	0.08	-0.14	0.53
Some WFH	0.06	0.70	-0.06

Notes: This table presents the factor loadings of a three-factor model using five variables in the EEC. Factor loadings are estimated using maximum likelihood. An oblimin rotation is used to allow for factors to be correlated rather than restricting them to be orthogonal.

Table 2 presents the factor loadings. The first factor loads positively on evening/night work and negatively on fixed schedules, capturing demanding and unpredictable hours. The second factor loads on occasional remote work, negatively on frequent remote work, and positively on working in the evenings. This factor capturing jobs where workers bring work home beyond regular hours rather than flexible hybrid arrangements. This factor is representative of "greedy work": how jobs place demands on workers' time outside of ordinary working hours. The third factor captures regular telework. In the estimation, I drop the third factor because preferences for working from home are likely to have changed dramatically since the COVID-19 pandemic.

2.2.3. Survey of Working Conditions and Psychosocial Risks (CT-RPS)

For descriptive analysis of amenity distributions, I supplement the EEC with the *Enquête Conditions de travail et risques psychosociaux* (CT-RPS), a detailed workplace conditions survey conducted every three years. The survey aims to understand workplace conditions and risks, and asks a detailed set of questions similar to the US O*NET survey. I use the 2016 and 2019 waves. Like the EEC, the CT-RPS includes occupation codes and establishment identifiers, allowing aggregation to the job

level (enterprise × occupation × location) for linkage to administrative data. The CT-RPS contains hundreds of workplace amenity questions. Following Goldin (2014), I reduce dimensionality using principal components analysis (PCA) to group correlated characteristics into summary measures. Table 7 presents the first eight principal components; detailed loadings are in Appendix D.2.

2.3. Sample selection and summary statistics

TABLE 3. Summary Statistics: Main Analysis Sample

	Mean
Panel A: Sample characteristics	
Female share	0.40
Mean age	40.05
Mean annualized earnings (€)	32,532.67
Permanent contract share	0.86
Median firm size	11
Number of workers	10,836,155
	Share of workers
Panel B: Occupational skill composition	
Executives + high-skill professionals	0.25
Intermediate professionals	0.20
Employees	0.23
Laborers	0.31
	Share of workers
Panel C: Sector Composition	
Construction	0.06
Diversified services	0.31
Industrials	0.20
Public administration	0.15
Trade, transport, and hospitality	0.27

Notes: This table provides summary statistics on the main analysis sample of the 2017-2019 and 2021-2023 BTS-Postes merged to the EEC. All means are computed as weighted averages, where weights are assigned using the full-time-equivalency variable in the BTS-Postes. Earnings variables are annualized using full-time equivalency to account for partial-year job spells. Mean annualized earnings are calculated by exponentiating the mean of the log annualized earnings distribution.

Because my model focuses on the worker's discrete job choice problem rather than on intensive margin labor supply, I restrict my main analysis sample to full-time workers. I also restrict to prime-age workers—25 to 55—and include only workers on permanent or fixed-term contracts (CDI

and CDD in the French classification), which excludes apprenticeships and temporary work. When constructing job-level aggregates, I weight individual job spells using the full-time equivalency reported in the administrative data. I also annualize earnings using the full-time-equivalent variable. This allows me to include job spells in my sample that are full-time but that start or end partway through the year.

Table 3 provides summary statistics on the sample of jobs that match to observations in the EEC data at the (enterprise, 2-digit occupation, département) level over the period of 2017-2023 (excluding 2020). In the Appendix I confirm that the EEC subsample is representative by comparing its characteristics to the full population meeting the same selection criteria.

2.4. Definition of labor markets using job-to-job transitions

To accurately estimate workers' labor supply elasticities in a discrete choice model, it is important to correctly specify the structure of their choice sets, or the *labor markets* in which they act. A labor market represents the reasonable alternatives available to a worker (Jarosch et al. 2024). It is clear that workers do not consider all jobs in the economy to be equally reasonable or equally substitutable for one another. But precisely how substitutable jobs are from the worker's perspective is an empirical question, and one whose answer informs us about how competitive labor markets are. Moreover, the sheer number of alternatives in administrative matched employer-employee data also poses a computational problem: estimation of a discrete choice model over so many alternatives is infeasible without imposing some additional structure.

To address these two problems, I develop a data-driven approach to identifying labor markets based on observed job mobility. This strategy follows the approach of Nimczik (2023) and Jarosch et al. (2024). I use panel data on job-to-job transitions to construct a network of jobs, where jobs serve as nodes and observed transitions between these jobs define directed, weighted edges. The network is then partitioned into distinct labor markets using a stochastic block model (SBM)—a standard approach in the network literature for detecting latent communities (Holland et al. 1983; Karrer and Newman 2011). In this framework, labor markets correspond to communities of firms that are tightly linked through mobility patterns. I provide more detail on the approach and the structural interpretation of the network graph in Appendix D.5.

The set of job-to-job transitions is measured over the Panel-BTS file period of 2013-2021. I define job-to-job transitions in the panel using the following restrictions. First, I include only transitions with no more than 30 days between the end of the first job and the start of the second job. Second, I require that workers hold full-time positions in each job and have tenure in each of at least one year. Third, I exclude transitions where either job is a temporary or apprenticeship position, though jobs can be either permanent or fixed-term contracts. Fourth, I require both jobs to be at

firms with at least five employees. Finally, I drop transitions for workers born before 1945 or after 1996 (who are either beyond retirement age or under the age of 25 in the whole panel) and workers for whom we only ever observe a single employment spell.

The SBM requires pre-specification of a fixed number of latent labor markets. I choose 20 markets to balance interpretability and goodness-of-fit. The fit of the partition to the observed network is captured by the modularity score, which measures the fraction of edges that fall within groups minus the expected fraction if edges were distributed at random. Modularity scores take values in $[-0.5, 1]$ and my chosen partition has a score of 0.94. This naturally aligns with the finding that 98 percent of job-to-job transitions in my network are within-market. I plot the modularity scores of networks with $k = 5, ..., 50$ communities in Appendix D.5.

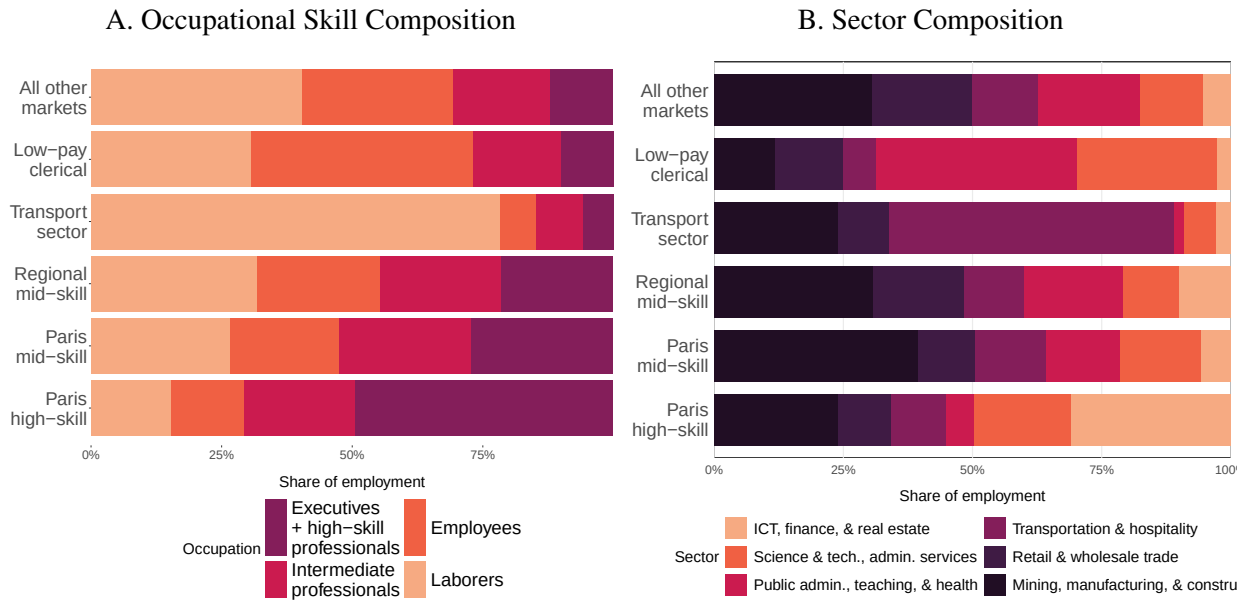
The markets capture observable differences across jobs in a variety of dimensions. There are five markets with clearly different compositions from the others—in terms of skill, sector, and/or geography—which I highlight as "major markets" in this section. Detailed summaries of skill, sector, and geographic compositions of all twenty markets are provided in Appendix D.5.

Figure 1 summarizes the occupational skill (Panel A) and sectoral composition (Panel B) of the five major markets. Panel A shows that the occupational skill distribution differs starkly across markets. The population of workers in the highest skill category, "Executives and high-skill professionals", is highly concentrated in the Paris high-skill market. The transport-sector market has a very high concentration of manual laborers. The low-pay clerical market has a disproportionate share of "employees", primarily consisting of administrative and clerical workers. All markets have at least some representation from every skill level, highlighting that occupation alone does not perfectly capture the boundaries of actual labor markets. Defining labor markets using occupational classifications would fail to capture connections between jobs in the data.

Panel B shows the sectoral composition of the major markets. The Paris high-skill market has a greater concentration in typically high-skill sectors—such as science and technology, information and communications, and finance and insurance—than other markets. However, the sectoral composition is diversified, with a significant share of workers in retail and wholesale trade and manufacturing as well. The mid-skill markets also represent a variety of sectors. Two markets show clear concentration in particular sectors: the transportation sector market concentrates nearly 60 percent of its workforce in transportation; workers in the low-pay clerical market are primarily employed in administrative services or in the public sector. Labor markets include alternatives across many sectors and caution against using sectors to define the boundaries of labor markets.

Table 4 summarizes key characteristics of the jobs in each of the major market groups, separately for men and women. The five identified labor markets exhibit distinct demographic and economic profiles that align with their occupational compositions. The Paris high-skill market stands out with

FIGURE 1. Occupational Skill and Sector Composition of Major Labor Markets



Notes: Panel A summarizes the share of jobs in each major labor market that belong to each 1-digit occupational group in the French PCS-ESE classification system, which broadly correspond to skill levels. Panel B summarizes the share that belong to each sector. Sector is defined as the French APEN code aggregated to 17 sector groups; I further aggregate into six categories. The lower five bars correspond to individual labor markets. The sector distribution for the remaining fifteen markets is aggregated in the top bar.

the highest annual earnings (€45,706 for men, €37,798 for women) and job stability (92.9 percent permanent contracts among men, 85.8 percent among women). It also represents the largest labor market in the economy, with 3.6 million workers predominantly in professional and managerial roles. Across markets, women earn 8-20 percent less than men in the full-time positions that form the sample. Women also work at slightly smaller firms. The table highlights that conditional on being in the same labor market, women hold compositionally different jobs from men and experience a substantial wage gap, even conditional on working full time.

3. Descriptive facts and stylized model

This section documents three facts about how firms set wages and amenities. First, most of the variation in amenities occurs at the firm level, with variation across occupations within the firm playing a secondary role. Second, higher-paying firms offer some better amenities (workplace safety, consistent schedules) but some worse amenities (more stress, evening work). Third, women sort into jobs with better amenities both across and within occupations. Together, these patterns motivate a model incorporating both worker preferences and firm technology constraints. Understanding the

TABLE 4. Gender Heterogeneity in Market Characteristics

	N	Female	Age		Earnings		Firm Size	
Labor market	**Workers**	**Share**	**Men**	**Women**	**Men**	**Women**	**Men**	**Women**
Paris high-skill	3,641,300	0.388	40.5	39.2	45,706	37,798	10	7
Paris mid-skill	662,916	0.395	41.3	39.6	38,949	30,638	9	6
Regional mid-skill	582,048	0.411	41.3	40.3	35,596	29,437	8	6
Transport sector	332,860	0.166	42.7	39.4	30,031	27,722	14	3
Low-pay clerical	1,014,420	0.469	41.0	39.5	26,370	23,861	9	9
All other markets	4,602,611	0.411	40.2	38.9	29,733	24,343	3	3

Notes: Labor markets are defined using a 20-market partition under a stochastic block model (SBM) following the approach of Nimczik (2023). Summary statistics are computed by individual markets and pooled for the "all other markets" category. Columns corresponding to age report mean values. The earnings value is calculated by exponentiating the mean value of log annualized earnings in each market, where earnings are annualized using the start and end date of the spell. All workers in the sample are full-time. Firm size statistic reflects the firm size of the median worker in the market. Partitions are computed using job-to-job transitions in the 2013-2021 file of the Panel-Base tous salariés.

roles of these parameters guides the estimation strategies in Sections 4 and 5.

Fact 1: Firms, then occupations within firms, explain most of the variation in amenities

The first fact explores how firms choose to provide different amenities. Much of the literature on amenities and compensating differentials takes amenities to be exogenous conditional on observable characteristics. This fact builds evidence that amenity provision is an endogenous firm-level strategic choice that is shaped by production technologies. Most of the variation in amenities occurs at the level of the firm, with variation across occupations within the firm playing a secondary role.

To do this, I decompose the sources of variation in amenities across jobs, similar to Hazell et al. (2022)'s decomposition of wage variation. For each amenity a, I regress the value of the amenity in occupation o at firm j in location l on combinations of firm, occupation, and location fixed effects:

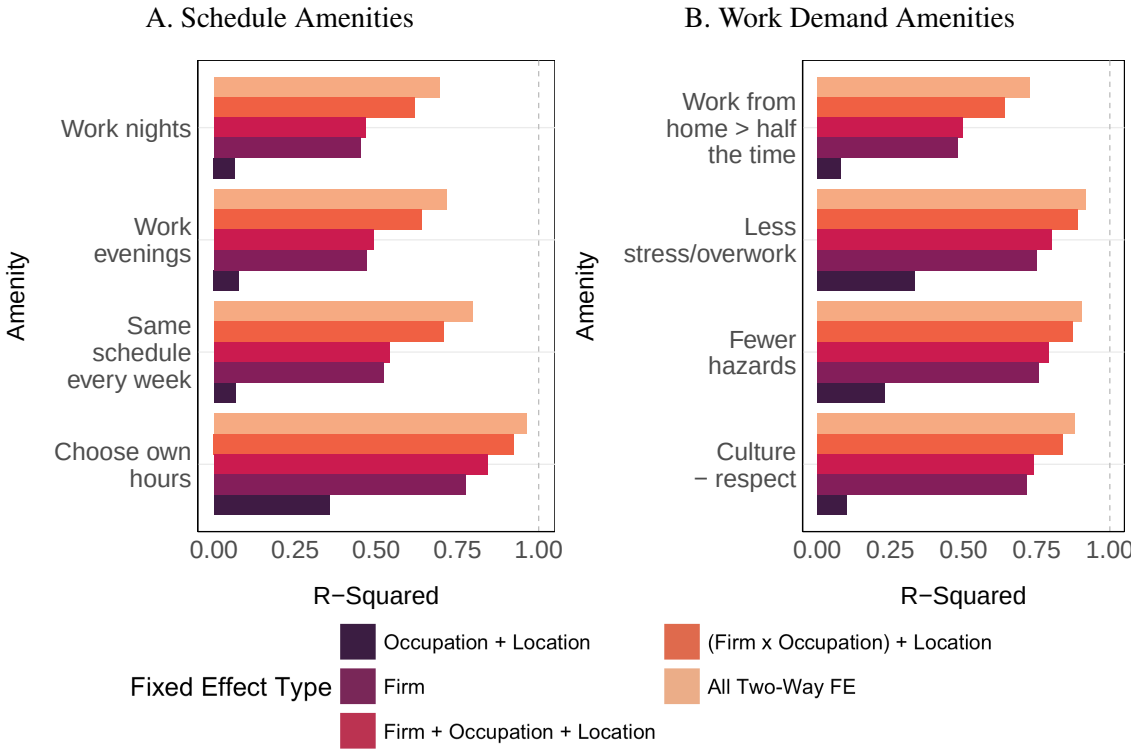
$$a_{ojl} = \gamma_o + \gamma_j + \gamma_l + \varphi_{oj} + \varphi_{jl} + \varphi_{ol} + \rho_{ojl} + \varepsilon_{ojl}, \tag{1}$$

where γ_o, γ_j, and γ_l represent one-way fixed effects for occupations, firms, and locations, respectively; φ_{oj}, φ_{jl}, and φ_{ol} represent the possible combinations of two-way interacted fixed effects; and ρ_{ojl} represents the three-way interacted fixed effect for each combination of firm-occupation-location.[5] The results are robust to the inclusion of year fixed effects.

Figure 2 illustrates how the adjusted R-squared of these regressions varies with the inclusion of

[5]Because I observe and link amenity data at the level of the enterprise × 2-digit occupation × département × year, including four-way interacted fixed effects will by construction explain 100 percent of the variation in the observed amenity a_{ojl}.

FIGURE 2. Amenity Characteristics Explained by Firms, Occupations, and Locations



Notes: This figure shows R-squared values from regressions of 1 with different types of interacted fixed effects (differentiated by bar color), with each amenity shown on the Y axis on the left hand side of its own regression. The following amenities are derived from the Continuous Labor Force Survey (EEC) from 2016 to 2023: Schedule varies week to week, Same schedule every week, Choose own hours. The following amenities are derived from the Survey of Working Conditions and Psychosocial Risks (CT-RPS) from 2016 and 2019: Schedule – more visibility/notice, Schedule – more stability. Amenities derived from the CT-RPS are measured as standardized values of principal components. Amenities are observed at the firm-occupation-location-year level, where firm is the SIREN (enterprise) code, occupation is the 2-digit French PCS-ESE code, and location is the département.

different combinations of fixed effects. Panel A focuses on amenities related to scheduling; Panel B show regression results for other amenities that capture demands that the job places on the worker. Across a variety of amenities, two common trends emerge. First, for all of the observed amenities, more than half of the variation is explained by firm fixed effects alone. For several amenities, the R-squared contribution of firm fixed effects alone is in the range of 70-75 percent. In contrast, including both occupation and location fixed effects does little to explain variation in amenities. Similarly, adding occupation and location fixed effects to the specification with firm fixed effects adds minimal explanatory power, increasing the R-squared value by less than 5 percentage points in each case.

Second, for many amenities, interacting the firm and occupation fixed effects adds an important secondary source of explanatory power to the regression. Consider the example of whether the

worker has the same schedule each week. 50 percent of the variation in this characteristic is explained by firm fixed effects alone. Interacting firm and occupation fixed effects raises the R-squared to 64 percent, suggesting that differences across occupations within the firm explain an incremental 14 percentage points of the variation in the outcome. The firm-level fixed effects are the most important source of variation in amenities, but occupational differences within firms playing a meaningful secondary role. I show R-squared decompositions for additional amenities in the Appendix. These additional decompositions illustrate that the same patterns of firmwide and occupation-specific amenity setting hold across a wide variety of job characteristics, including firm culture, additional measures of flexibility, and the frequency of work during odd hours. These results suggest that firms make coordinated decisions about amenity provision across jobs.

Fact 2: Better-paying firms offer some better and some worse amenities, providing evidence of *both* compensating and augmenting differentials

Do higher-paying jobs offer better or worse amenities? This question relates to a long-standing debate in the literature over whether the value of amenities amplifies or offsets pay inequality in the labor market. On the one hand, better firms may offer better amenities, leading amenity provision to exacerbate inequality across workers. On the other hand, amenities may act as compensating differentials—so that workers give up wages in exchange for valued amenities—which would imply that they serve to dampen inequality.

I answer this question by correlating observed amenities with job-specific wage premia following the definition of Bonhomme et al. (2019), henceforth BLM. Within each sector, jobs are classified into ten categories using k-means clustering. Following BLM, I cluster jobs on 20 quadrature points of the job-specific wage distribution in the 2018 cross-sectional BTS-Postes data set. The wage is defined as the full-time-equivalent annualized wage and the job is defined at the level of the enterprise × two-digit occupation × region. The full-time-equivalence conversion variable in the BTS-Postes allows me to adjust observed earnings both for spells that are shorter than the full year and for less-than-full-time work. Clusters are ranked according to the mean full-time-equivalent wage of firms that are assigned to the cluster. I describe the clustering approach in greater detail in Appendix D.3.

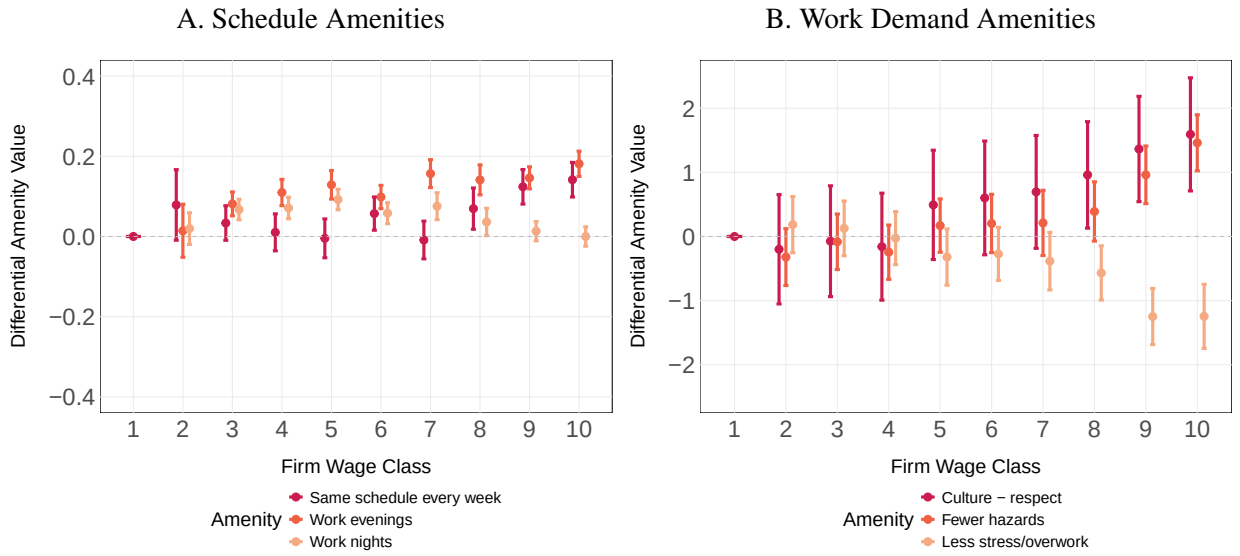
For each amenity a, I estimate:

$$a_{ojlt} = \boldsymbol{\alpha} \cdot \boldsymbol{D_j^{BLM}} + \gamma_t + \varepsilon_{ojlt}, \tag{2}$$

where D_j^{BLM} is a vector of $k = 1, ..., 10$ indicators that job j is a member of BLM wage class k and γ_t is a vector of year fixed effects. In the specification with year fixed effects there is no intercept

and the coefficient α_1 corresponding to the lowest-ranked BLM wage class is unidentified. Because BLM wage classes are estimated within sectors, this specification implicitly controls for systematic differences in amenities in low-paying versus high-paying sectors.

FIGURE 3. Different Amenities are Correlated with Firm Wages in Different Ways



Notes: This figure shows point estimates of the coefficients α_k on each BLM wage class $k = 2, ..., 10$ in Equation 2. Estimates represent the differential amenity value in standard deviations relative to the class of firms with the lowest wages. Regressions include year fixed effects. Amenity measures are derived from the Survey of Working Conditions and Psychosocial Risks (CT-RPS) from 2016 and 2019. Amenities derived from the CT-RPS are measured as standardized values of principal components. Amenities are observed at the firm-occupation-location-year level, where firm is the SIREN (enterprise) code, occupation is the 2-digit French PCS-ESE code, and location is the département. Standard errors are clustered at the SIREN (enterprise) level.

Figure 3 plots estimates of α_k for $k = 2, ..., 10$. Panel A shows results for amenities related to scheduling measured in the EEC. Panel B shows results for amenities that capture other demands that the job places on the worker, measured in the CT-RPS. Coefficients in Panel A represent differences in the share reporting evening or night work relative to the lowest-paid category. Coefficients in Panel B represent standard deviation differences in amenity values relative to the lowest-paying set of jobs.

Higher-paying jobs have distinct bundles of amenity characteristics related to scheduling–some better, some worse. Panel A shows workers in the highest pay class work evenings 14.5 percentage points more often than those in the lowest class, while middle classes report evening work 5-9 percentage points more often than those in the lowest class. Night work shows a non-monotonic pattern: firms in pay classes 3-7 are significantly more likely to require it, but the highest-paying firms are not. The highest-paying jobs are also 10-15 percentage points more likely to offer consistent weekly schedules. These patterns suggest that amenity differences cannot be explained by a simple

one-dimensional trade-off against wages.

Panel B shows how workplace demands–physical, psychological, and cultural–vary across jobs that pay better or worse. High-paying jobs are 1.46 standard deviations less likely to expose workers to workplace hazards, such as carrying heavy objects or handling hazardous materials. Since pay classes are computed within a sector, this pattern cannot be explained by sectoral differences in hazard exposure that are also correlated with pay. Workers in the highest-paying jobs are also 1.59 standard deviations more likely to report that they are treated with respect than workers in the lowest-paying jobs in the same sector. However, high-paying jobs also have substantially worse stress and overwork, with the highest-paying class scoring 1.22 standard deviations worse. These results suggest roles for both compensating and augmenting differentials in the provision of amenities, depending on the amenity type. While high-paying jobs have lower exposure to workplace hazards and cultures where workers feel more respected, they often come with greater stress and overwork.

Fact 3: Women tend to sort into jobs with better amenities than men, mostly driven by heterogeneity *within* occupation rather than by selection *across* occupations

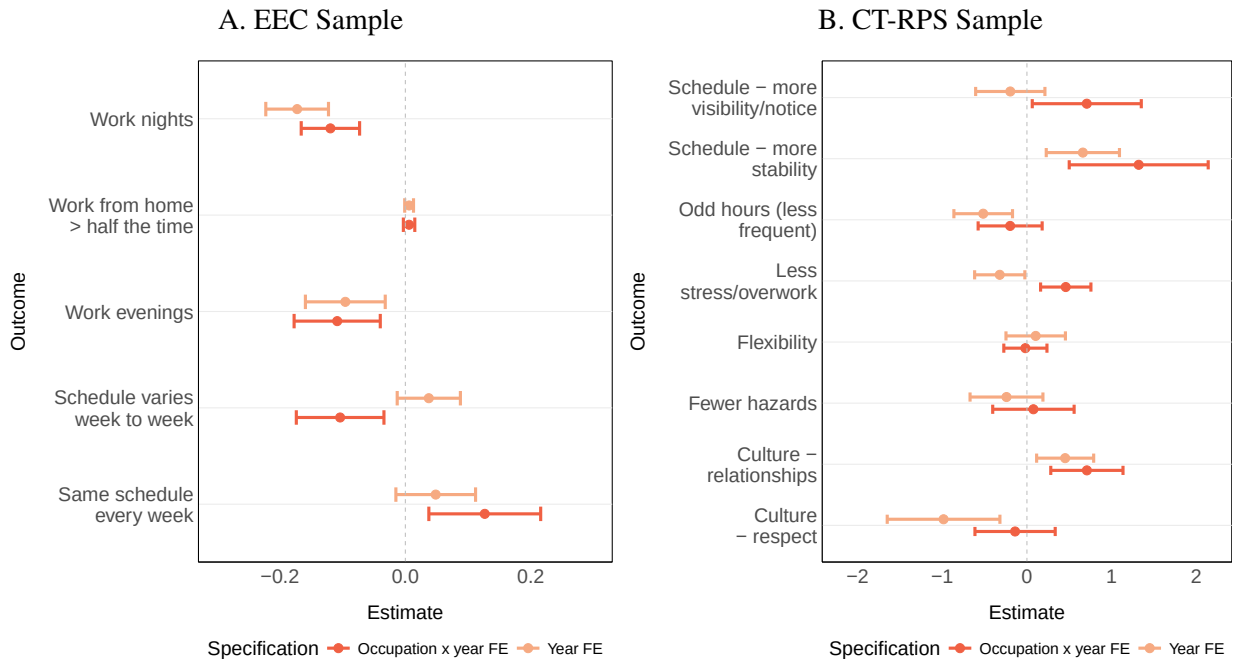
I now ask: do men and women sort systematically to firms that offer different amenities? This matters for understanding the role of supply-side preferences in observed amenity variation. I document that women's jobs tend to have better amenities—particularly in relation to scheduling—in a way that is mostly explained by sorting within rather than across occupations. To establish this fact, I regress the amenity value of job ojl, a_{ojlt}, on the share of workers in that job who are female, f_{ojlt}, as well as fixed effects for year γ_t or for occupation-year γ_{ot}, depending on the specification:

$$a_{ojlt} = \alpha + \beta \cdot f_{ojlt} + \gamma_t + \gamma_{ot} + \varepsilon_{ojlt}. \tag{3}$$

In my preferred specification, occupation is defined as 2-digit occupation × sector to allow for sectoral differences in occupational amenities. Intuitively, the regression recovers the average difference in the level of amenity a for women relative to men, alternatively including and excluding controls for occupational sorting. Standard errors are clustered at the enterprise level and regressions are weighted by the number of workers in each job.

Figure 4 plots the point estimates for coefficients β in these regressions for each observed amenity. Panel (a) plots point estimates for the EEC sample, which can be interpreted as the percentage point increase in share of workers who report having the (dis-)amenity associated with a 1-percentage-point increase in the female share in that job. Panel (b) plots point estimates for the CT-RPS sample, which can be interpreted as the standard deviation increase in the quality of the amenity associated with an all-female job relative to an all-male job.

FIGURE 4. Gender Differences in Job Amenities



Notes: This figure shows point estimates of the coefficient β from regressions of Equation 3 with different types of interacted fixed effects (differentiated by bar color), with each amenity shown on the Y axis on the left hand side of its own regression. Amenities derived from the EEC are indicators that the job has the characteristic indicated. Amenities derived from the CT-RPS are measured as standardized values of principal components. Amenities are observed at the firm-occupation-location-year level, where firm is the SIREN (enterprise) code, occupation is the 2-digit French PCS-ESE code interacted with 10-category NAF sector code, and location is the département.

Women's jobs have worse amenities on two dimensions—sense of respect and stress/overwork—but these differences are driven by occupational sorting. Once I control for occupation-year fixed effects, gender differences in respect disappear, and women report *lower* stress and overwork. All other amenity differences persist after controlling for occupation, indicating women sort to better amenities primarily *within* rather than across occupations. Women have jobs with greater schedule visibility and stability, better workplace relationships, and lower rates of evening and night work. Women also sort into jobs with frequent remote work at slightly higher rates, though the small estimate masks substantial variation over time given the recent emergence of widespread remote work.

3.1. A Stylized Model of Firms' Wage and Amenity Setting

The facts raise three questions. First, why do amenities co-move within firms (Fact 1)? Second, why are some amenities positively correlated and others negatively correlated with pay (Fact 2)? Third, how much of gender sorting reflects worker preferences versus firm technology (Fact 3)? To organize thinking about these questions, I develop a stylized model of how firms choose wages

and amenities. The model rationalizes the observed facts and clarifies which economic primitives govern equilibrium compensation packages, motivating the estimation approach in Sections 4 and 5 as well as the counterfactuals in 6.

In the stylized model, firms firm choose occupation-specific wages and a flexibility amenity to maximize output net of labor costs. The firm uses two types of labor inputs, $o \in \{L, H\}$, and chooses wages and a single amenity to offer to each, w_{ojt} and a_{ojt}, respectively. I assume in the stylized model that workers are identical in their unobserved ability. The amenity represents the degree of flexibility afforded to the worker. Flexibility may reduce worker productivity through a job-specific shifter z_{ojt}, which depends on the *flexibility cost* captured by the parameter $\delta_{oj} < 0$. Production follows a CES technology over occupation-specific effective labor inputs:

$$Y_{jt} = A_{jt} \cdot \left((z_{Hjt} L_{Hjt})^{\rho_j} + (z_{Ljt} L_{Ljt})^{\rho_j} \right)^{\frac{\alpha_j}{\rho_j}}, \quad \text{where} \quad z_{ojt} = \exp(\delta_{oj} a_{ojt}). \tag{4}$$

Firms choose wages and amenities for each occupation $\{w_{ojt}, a_{ojt}\}_{o \in \{L,H\}}$ to maximize profits, taking labor supply as a function of wages and amenities. I develop a model of labor supply in greater detail in Section 4; for the stylized model, we only need the firm to face an upward-sloping labor supply curve so that labor supply is imperfectly elastic. The firm then maximizes profit subject to the constraint that labor supply equals labor demand.

I consider an equilibrium in which each firm (and each job) is strategically small within the labor market. This implies that firms ignore the effect of changing the wage for the job in occupation H on labor supply to the job in occupation L. In Appendix C.1, I derive the general first-order conditions allowing for cross-occupation effects—that is, allowing the wage or amenity in job Hj to affect labor supply to job Lj, and vice versa. I show in the Appendix that the equilibrium expressions below obtain when firms are strategically small and cross-occupation elasticities are negligible. This assumption can be rationalized by the facts that (i) individual jobs have small market shares, and (ii) jobs in different occupations are not close substitutes in workers' choice sets. Point (i) is easily verified by examining the empirical market shares of individual jobs; point (ii) holds in the labor supply system estimated in Section 4. The assumption also aligns closely with related work in this literature (Lamadon et al. 2022; Kroft et al. 2025).[6]

Firms choose optimal wages as a markdown to the marginal product of labor. The markdown is a function of the elasticity of labor supply with respect to the wage, which is derived from the labor

[6]While related papers in this literature do not account for multiple jobs within the same firm, my assumption is consistent with theirs so long as a worker do not view a job that is within her current firm as more substitutable than an otherwise identical job outside the current firm. I make this assumption in the substitution structure of my labor supply model. Moreover, recent papers testing for strategic interactions in wage setting (Sharma, 2023; Roussille and Scuderi, 2024; Derenoncourt and Weil, 2025) do not find evidence of such interactions, suggesting that a model of monopsonistic competition as described here better captures firm behavior.

supply curve. Taking first order conditions leads to the equilibrium wage, expressed in logs:

$$\log w^*_{ojt} = \underbrace{\log\left(\frac{\varepsilon^w_{ojt}}{1+\varepsilon^w_{ojt}}\right)}_{\text{markdown}} + \underbrace{\log\alpha_j + \log A_{jt} + \log\left(\frac{Y_{jt}}{A_{jt}\cdot L_{jt}}\right)}_{\text{firm productivity}(\psi_{jt})} + \underbrace{\rho_j\delta_{oj}a_{ojt}}_{\text{amenity cost}} + \underbrace{(\rho_j-1)\log L_{ojt}}_{\text{returns to scale}}, \quad (5)$$

where ε^w_{ojt} is the wage elasticity of labor supply and $L_{jt} = (z_{Hjt}L_{Hjt})^{\rho_j} + (z_{Ljt}L_{Ljt})^{\rho_j}$. This expression clarifies that wages depend on four forces: monopsony power (the markdown); firm productivity ψ_{jt}, including both returns to scale in labor and firmwide productivity shocks; the direct cost of providing amenities ($\rho_j\delta_{oj}a_{ojt}$); and within-occupation returns to scale. Crucially, the relative importance of each of these determinants of the wage depends on economic primitives: wage elasticities (ε^w_{ojt}) and firm technology parameters (α_j, ρ_j, and δ_{oj}).

The level of the flexibility amenity is chosen to attract more labor at a lower wage, taking into account that providing the amenity may incur a cost by reducing workers' productivity. The firm's optimal amenity provision equates marginal returns and marginal labor supply costs. Expressed in logs, the equilibrium amenity is:

$$\log a^*_{ojt} = \underbrace{\log\left(\frac{\varepsilon^a_{ojt}}{1+\varepsilon^w_{ojt}}\right)}_{\substack{\text{relative amenity}\\ \text{preference}}} - \underbrace{\log(-\delta_{oj})}_{\text{amenity cost}}, \quad (6)$$

where ε^a_{ojt} is the elasticity of labor supply with respect to the amenity. Firms will provide more flexibility when workers are more sensitive to the amenity *relative* to wages (high $\varepsilon^a/\varepsilon^w$) and when flexibility is less costly to provide (low $|\delta_{oj}|$). This highlights two distinct channels for amenity setting that I will separately identify through estimation of supply and demand. Full derivations of the equilibrium conditions for wages and amenities are provided in Appendix C.1.

3.2. Interpreting the facts through the model

The model rationalizes the three facts by clarifying how worker preferences and firm production parameters jointly determine equilibrium amenities.

Why do amenities co-move within firms? The model predicts that amenities will be positively corrleated across jobs within a firm when either of two forces line up across occupations: (i) *labor market conditions* (similar wage and amenity elasticities), or (ii) *costs of provision* (similar technologies to provide flexibility). Equation 6 shows that amenities depend on labor supply parameters (ε^a, ε^w) and the cost to provide amenities. The secondary role of occupation within the

firm reflects flexibility costs and labor market conditions that may vary across different kinds of jobs within the firm.

What determines the sign of the wage-amenity correlation? The sign depends on the covariance between firm productivity (A_{jt}) and flexibility costs (δ_{oj}). When high-productivity firms are efficient at providing amenities, both wages and amenities are high. The data show that high-paying jobs tend to have fewer physical hazards, more consistent schedules, and a greater sense of respect in the workplace—suggesting that more productive firms do a better job of providing these attributes. When amenities penalize output, wages are high but amenities are low (compensating differentials). The data are consistent with this phenomenon as well: higher-paying jobs entail more stress and evening work, which are tightly linked to worker productivity.

Why do women sort to jobs with better amenities within occupations? If women value amenities more strongly relative to wages than men do, they will disproportionately sort to firms offering better flexibility. From the firm's perspective (equation 6), attracting more amenity-sensitive workers makes it optimal to provide more amenities, reinforcing the sorting pattern.

3.3. What will we learn from estimation?

The model clarifies the key economic parameters that I will estimate in order to quantify the contributions of labor supply and labor demand to amenity heterogeneity, gender sorting, and differences in pay. First, the job-specific elasticities of labor supply with respect to wages and amenities measure the firm's labor market power and workers' sensitivity to amenity provision. These parameters determine the extent to which firms choose amenities in response to labor market conditions. In Section 4, I estimate a labor supply model with preference heterogeneity and endogenous amenities to recover these elasticities for individual jobs. Second, the productivity costs faced by firms providing amenities, δ_{oj}, measure the role of technological constraints in firms' equilibrium choices. In Section 5, I estimate a model of labor demand to recover occupation-specific flexibility costs. Understanding the importance of these two channels will teach us about how mandating amenities affects worker welfare, taking into account firms' equilibrium choices and workers' endogenous sorting.

4. Labor supply: how do workers value wages and amenities?

4.1. Theoretical framework and estimation approach

In this section, I describe how workers with heterogeneous preferences choose a job to maximize utility. In the model, I define a *job j* as an occupation *o* performed at firm *k* in location *l*. For example,

all software engineers at Firm A in Paris hold the same job; software engineers at the same firm in the Hauts-de-Seine region hold a different job. Each job belongs to a unique labor market $m \in \mathcal{M}$ and is characterized by a wage w_{jt}, a vector of amenities $\boldsymbol{a_{jt}}$, and a vector of exogenous amenities $\boldsymbol{y_{jt}}$. Wages and endogenous amenities are equilibrium objects that are chosen by firms, reflecting both firm technology parameters and labor supply elasticities as discussed in Section 3.1. Exogenous amenities can include features such as the commuting distance between the worker and the firm that are plausibly exogenous to shocks to worker preferences in period t. In this exposition, I suppress time subscripts t to reduce notation since the problem is static. I summarize the job j using the tuple $j = (k, m, o, l)$.

A worker i with demographic characteristics $d(i)$ and ability X derives utility from a job $j = (k, m, o, l)$ through four inputs: (1) the wage $w_j(X)$, which can depend on her ability; (2) endogenous amenities $\boldsymbol{a_j}$; (3) exogenous amenities $\boldsymbol{y_{ij}}$, which can include characteristics that are specific to the worker-job match (here, I study commuting distance); (4) unobserved job quality $\xi_{d(i)j}$; and (5) i's idiosyncratic preference for the alternative (j, o, l), captured by the error term ϵ_{ij}. The utility of the worker can be written as

$$u_{ij} = \underbrace{\alpha_{d(i)} \log w_j(X)}_{\text{wage value}} + \underbrace{\boldsymbol{a_j'\beta_{d(i)}}}_{\substack{\text{endogenous} \\ \text{amenity value}}} + \underbrace{\boldsymbol{y_{ij}'\psi_{d(i)}}}_{\substack{\text{exogenous} \\ \text{amenity value}}} + \underbrace{\xi_{d(i)j}}_{\substack{\text{unobserved} \\ \text{job quality}}} + \underbrace{\epsilon_{ij}}_{\substack{\text{idiosyncratic} \\ \text{preference}}} . \tag{7}$$

The specification of the worker's utility function is similar to the model of Fan (2013) for consumers' preferences over products with endogenous characteristics. I describe each of the components of the utility function in greater detail below.

$w_j(X)$ is the equilibrium wage received by the worker of ability X at job j. As derived in the stylized model of section 3.1 and shown in equation 5, the equilibrium wage reflects a markdown to the marginal product of the worker. Note that wages are allowed to vary across jobs j within the same firm k and across ability types X within the job, but cannot vary within the job by demographics d. This reflects non-discrimination regulations in France that are common in developed economies: workers must receive equal pay for equal work, but more productive workers may be paid more than less productive workers.

Jobs provide a vector of endogenously chosen amenities $\boldsymbol{a}$ that are specific to the job j. Like wages, amenities cannot vary by demographic group, but can differ across jobs within the firm. The value of a given level of the mth amenity a_j^m is determined by the wage-equivalent value β_{md}, which characterizes the relative importance of a^m in the utility function for worker i of demographic group $d(i)$. The worker's valuations of amenities, both endogenous and exogenous, are assumed not to vary across jobs. Demographic heterogeneity in preferences for the mth endogenous amenity is

captured by

$$\beta_{md} = \beta_m + \theta^{\beta}_{md}.$$

β_m captures the common value of the quality of the amenity, on which all workers agree. θ^{β}_{m} is a vector of demographic group fixed effects that captures variation in amenity preferences across those demographic groups. For example, women may value flexibility or remote work more strongly than men (Mas and Pallais, 2017). For simplicity, I currently assume that preference parameters for amenities do not vary by ability type X.[7] Thus the only way that the ability X enters utility is through the equilibrium wage. I also allow for workers of different abilities and demographics to face different choice sets, and allow them to have different substitution elasticities across jobs.

Jobs can also have observable characteristics that are taken to be exogenous in the model because they are stable over time and do not respond to worker preferences in a static setting. This can include firm location (urban or suburban), whether the location is the headquarters of the firm, and the sector of the firm. Exogenous characteristics can also be specific to the worker-job match, such as the commuting distance between a particular worker and job. This feature can be considered exogenous if the firm's location in the static setting does not respond to the locations of individual workers. As with the endogenous characteristics, demographic heterogeneity in preferences for the mth exogenous amenity is captured by a coefficient shifter:

$$\psi_{md} = \psi_m + \theta^{\psi}_{md}.$$

The coefficient structure of the exogenous amenities mirrors that of the endogenous amenities. ψ_m captures the common value of the quality of the amenity, on which all workers agree. θ^{ψ}_{m} is a vector of demographic group fixed effects that captures variation in amenity preferences across those demographic groups. For example, women may have a stronger distaste for long commutes than men do (Le Barbanchon et al., 2020). Again, I assume that workers of different abilities X do not vary in their preferences for the exogenous amenities.

Finally, ξ_{djX} is the unobservable market-year-specific taste of demographic d for job j, which captures unobserved differences in quality across jobs that are not explained by observable wages or amenities. ξ_{dj} can include any non-pay characteristics of the firm that are valued by workers but are unobserved to the econometrician, such as the firm's reputation or workplace culture. Including the ξ_{dj} term allows us to learn about the amenities I can observe without requiring that they exhaustively

[7] A generalized extension of the model in which I estimate the entire labor supply system separately for workers of different types is currently in progress. In this version of the model, I allow workers of different types to have different preferences for wages and amenities.

capture all of the job's non-pay features. Finally, ϵ_{ij} is a Generalized Extreme Value preference shock that captures idiosyncratic preferences worker i may have for job j beyond what can be explained by the data.

I can summarize the mean value of a job in terms of McFadden's "alternative-specific constant" μ_{dj} corresponding to each demographic group. These objects are given by

$$\mu_{djX} = \alpha_d \log w_j(X) + \boldsymbol{a}_j' \boldsymbol{\beta}_{\boldsymbol{d}} + \boldsymbol{y}_j' \boldsymbol{\psi}_{\boldsymbol{d}} + \xi_{djX}. \tag{8}$$

Next, to obtain an equation for the probability that a worker with demographics d and ability X chooses a given job j, I develop a structure of the worker's consideration set that captures observed substitution patterns between jobs and motivates my estimation procedure.

Structure of the worker's consideration set and estimation procedure

To estimate the discrete choice model outlined above, I develop a two-stage approach adapted from the industrial organization literature on product market demand. My strategy builds on the framework of Azar et al. (2022), who use a nested logit choice model to estimate wage elasticities. In the first stage, I estimate job-specific mean utilities and substitution elasticities between alternatives. In the second stage, I estimate the contribution of wages and amenities to these mean utilities, which in turn allows me to recover job-specific labor supply elasticities with respect to both wages and amenities.

A key challenge in implementing this strategy lies in defining the set of alternatives over which workers choose. In practice, workers do not consider all jobs in the economy to be equally substitutable. Instead, each worker's consideration set is shaped by a combination of occupational specialization, geographic constraints, and labor market segmentation. I capture the structure of job choice using a nested logit framework, which partitions the universe of job alternatives into economically meaningful groups.

The nesting structure captures three levels of the worker's job choice problem. At the top level, I take as given that the worker is assigned to a **market**.[8] Each job in the economy belongs to a unique labor market $m \in \mathcal{M}$. Markets are constructed nonparametrically using job-to-job transition data following the approach of Nimczik (2023). The approach relies on the Stochastic Block Model

[8]One may also hypothesize that the worker chooses the market at a third, top level of the nested logit structure rather than taking the market assignment as given. In a previous version of this paper, I derive a three-tiered version of the nested logit and estimate parameters of this candidate model. I find that the top-level parameter Λ governing substitution across markets is precisely estimated and statistically indistinguishable from zero. Moreover, the resulting estimates for the λ_m substitution parameters in each market are outside the expected range of $(0, 1)$ under this alternative structure. I argue that this combination of results rejects the three-tiered structure in favor of a two-tiered structure where the market assignment is taken as given.

(SBM) (Holland, Laskey, and Leinhardt, 1983; Karrer and Newman, 2011), which is used in the network analysis literature for the detection of latent communities in network data. Intuitively, two jobs are in the same market if they are connected by frequent job-to-job transitions. Rather than relying on pre-defined groupings, this approach infers the latent labor market structure from the observed flow of workers and captures *revealed* substitutability between alternatives. This approach allows for flexible market boundaries that reflect actual worker flows rather than imposing arbitrary geographic or sectoral partitions. I provide greater technical detail on the construction of these labor markets in Appendix D.5. I select a partition of the economy into 20 discrete labor markets. Under the 20-market partition, 98 percent of employment-to-employment transitions occur within markets. This suggests that the market partition is consistent with observed patterns of worker mobility.

At the second level of the nesting structure, the worker chooses an **occupation conditional on the market** $o \in \mathcal{O}_m$. Occupations are defined at the 2-digit level using the French PCS-ESE occupation coding scheme. This secondary level of the nesting structure reflects the hypothesis that workers are more likely to substitute between jobs with similar functions and skill content. Estimating the substitution elasticity across occupations within the same market allows me to empirically test this hypothesis.

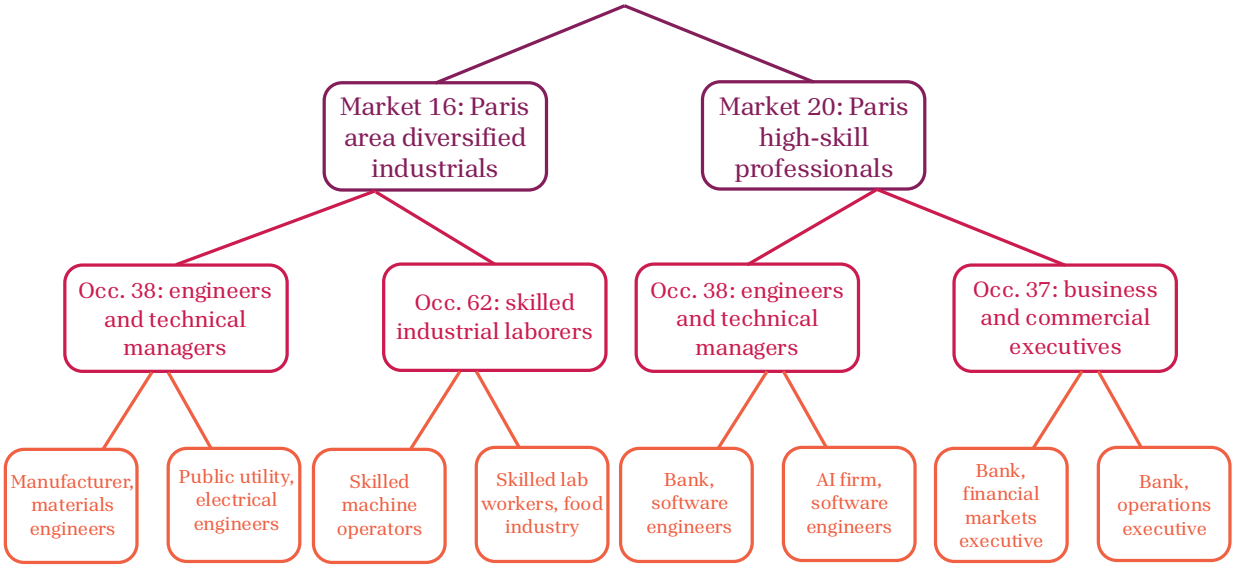
Last, conditional on the market m and occupation o, the worker chooses a **job** $j \in \mathcal{J}_{mo}$. Each job is defined by a combination of an enterprise (SIREN in the French administrative data), 2-digit occupation, and location (*département*, or region, in France). Formally, the nested logit structure requires that the Independence of Irrelevant Alternatives (IIA) assumption holds within, but not across, nests. Under this assumption, adding a new job to market-occupation nest (m, o) does not cause proportional substitution from all jobs in the economy—rather, it causes *disproportionately strong* substitution away from other alternatives within (m, o), with proportional substitution patterns for all jobs within (m, o). This assumption can be evaluated empirically by testing the hypotheses of and $\lambda_m = 1$ for each m. Substitution elasticity estimates between 0 and 1 imply the nesting structure is correctly specified. Importantly, the estimation routine does not mechanically guarantee values of $\lambda_m \in (0, 1)$, so I can empirically test whether the structure is well-specified.

There are both economic and computational motivations for the nesting structure. Economically, workers' substitution patterns reflect nested decisions: workers tend to first consider connected labor markets, then occupations within those markets, and finally specific firms. Empirically, 98 percent of job-to-job transitions are within-market: conditional on the choice of market, workers are dramatically more likely to move to another job within the same market than to change markets. Conditional on market, 70 percent of job-to-job transitions are within the same occupation.[9]

[9]This statistic excludes E:E transitions within the same firm. Including within-firm transitions, 29 percent of transitions are within the same occupation. However, within-firm job are much more likely to reflect promotions. In practice, including or excluding these does not change the structure of the labor market clusters.

Capturing these structured substitution patterns improves the plausibility of the estimated elasticities and avoids overstating the substitutability between dissimilar jobs. Computationally, the number of job alternatives in the economy in a given year is extremely large, making the estimation of job-specific mean utilities for each job in a flat logit model infeasible. The nested framework allows me to estimate the model sequentially, improving tractability without sacrificing rich substitution patterns across more or less similar alternatives.

FIGURE 5. Example Nesting Structure of Discrete Choice of Labor Supply



Notes: Illustrative nesting structure uses generic examples, not reflective of actual firms in the administrative data. The nesting structure depicted represents four levels of nested choices: (1) market, defined using the methodology of Nimczik (2023); (2) occupation, defined using the 2-digit French PCS-ESE standard occupation code; and (4) job, defined as a specific firm and location within a particular market-occupation nest. Firms are defined using 9-digit French SIREN codes, a unique enterprise identifier. Further details on the construction of markets at the first level of nesting are provided in Appendix D.5.

Figure 5 provides illustrative examples of the choice alternatives at each level of the nesting structure. At the top level, workers are assigned to one of many available markets in the economy, which vary widely in their occupational and geographic composition. This example shows two real markets with relatively high-skill occupational compositions: (1) a cluster of jobs in diversified industrial sectors–such as retail trade, manufacturing, and construction–concentrated around Paris, and (2) a cluster of high-skill executive and white-collar jobs in the Paris area. In this version of the model I assume that workers do not choose their market. In the Appendix I test and reject an alternative structure for the job choice problem in which workers choose among markets at the top

level of nesting.

Conditional on the market, workers choose a 2-digit occupation. The occupation alternatives in each market highlight the importance of occupation for grouping jobs into nests. For example, in the Paris professional cluster, two common occupational alternatives partition the choice into engineering and technical management jobs–common in technology firms–and business and commercial executives–who may work in similar firms, but have commercial, rather than technical, skill sets. In contrast, the occupational composition of the diversified industrial market highlights how different groups of occupations can be strongly connected in different labor markets. Engineers and technical managers are an important part of this labor market as well, but the next-most-common occupation is skilled industrial laborers. This example highlights that data-driven labor markets capture nuances in job groupings that segmenting labor markets *ad hoc* by location and occupation can miss. If I defined "labor markets" by location and occupation instead of by job-to-job transitions, engineering and technical managerial jobs in these two markets would be erroneously grouped together, misrepresenting the true choice set of workers in each.

Finally, at the bottom level, workers choose between jobs—defined by firm and location—within a market and occupation. For example, some of the options available to engineers and technical managers in the industrial cluster may include materials or electrical engineering jobs. In contrast, workers in the high-skill white-collar cluster face a different set of alternatives, such as software engineering positions at a bank or an AI firm. Of course, in the data, workers face many more alternatives within each nest.

Having defined the three-tiered nesting structure, I now discuss the two estimation stages. In the first stage, I sequentially estimate the nested logit model to recover job-specific mean utilities and elasticities of substitution between occupations within each market. In the second stage, I develop an instrumental variables strategy to estimate the effects of wages and amenities on job mean utilities and recover job-specific labor supply elasticities.

Estimation of bottom-level logit: choice of job conditional on market and occupation

The first stage of the nested logit estimation strategy proceeds in three steps, beginning with the estimation of the job choice logit at the bottom level of the nesting structure. This part recovers relative job utilities and commuting disutility parameters up to a normalization that is specific to the market-occupation nest. All parameters are estimated via maximum likelihood.

As before, define a *job* as a firm × occupation × region. Index markets by m and occupations by o, and jobs by j. The utility of worker i, who has demographic type $d(i)$ and ability $X(i)$, from job j is

$$u_{ijt} = \mu_{d(i)jX(i)} + \gamma_{mo} \cdot z_{ij} + \nu_{imo}(\lambda_m) + \lambda_m \epsilon_{ij}.$$

z_{ij} is the log commuting distance in kilometers that i would face by working at job j. The commuting distance is calculated using the latitude and longitude of the *communes* (municipalities) of the individual's residence and of the job.[10]

$\nu_{imo}(\lambda_{mdX})$ and ϵ_{ij} are random utility shocks that have the necessary distributions to generate the nested logit error structure as in Cardell (1997) and Azar et al. (2022). ϵ_{ij} represents an idiosyncratic match effect between worker i and job j within the nest. λ_{mdX} is a nesting parameter that governs the substitutability of occupations o within a market nest. λ_{mdX} is specific to the demographic group and ability type.

The probability that individual i chooses job j conditional on the most granular nest – i.e., on having already chosen occupation o within market m, and thus facing alternatives $j \in \mathcal{J}_{mo}$ – is given by

$$s_{ij|m,o} = \frac{\exp[(\mu_{d(i)jX(i)} + \gamma_{mo} \cdot z_{ij})/\lambda_{md(i)X(i)}]}{\sum_{j' \in \mathcal{J}_{mo}} \exp[(\mu_{d(i)j'X(i)} + \gamma_{mo} \cdot z_{ij'})/\lambda_{md(i)X(i)}]}. \tag{9}$$

Note that only the *relative* demographic-specific mean utilities ($\tilde{\mu}_{djX} \equiv \mu_{djX} - \mu_{dr(m,o)X}$) are identified in the bottom-level logit estimation, where $r(m,o)$ is a reference job specific to the market-occupation nest. To emphasize this point, we can rewrite the choice probability

$$s_{ij|m,o} = \frac{\exp[(\tilde{\mu}_{jd(i)X(i)} + \gamma_{mo} \cdot z_{ij})/\lambda_{mdX}]}{\sum_{j' \in \mathcal{J}_{mo}} \exp[(\tilde{\mu}_{j'd(i)X(i)} + \gamma_{mo} \cdot z_{ij'})/\lambda_{mdX}]}. \tag{10}$$

Maximum likelihood estimation in this step thus recovers the relative mean utilities $\tilde{\mu}_{djX}$ and commuting disutililty parameters γ_{mo} up to scaling by the substitution parameter λ_{mdX}. The reference demographic-specific utilities $\mu_{dr(mo)X}$, and nesting parameters λ_m will be identified in subsequent estimation of higher-level logits, allowing us to compare the utility values of alternatives across nests. I elaborate on how these nest parameters are recovered in section 4.1.

Estimation at this stage also requires the choice of a reference job or "outside option" in each nest. I construct a composite "reference job" by aggregating all jobs within the nest that are held by fewer than 25 employees into a single alternative, representing 2-5 percent of employment in each granular labor market. All characteristics of jobs in the outside option, including mean log annualized wages and amenities, are taken to be weighted averages of the characteristics of the constituent jobs. Weights are assigned according to the number of workers observed in that job in

[10]Because jobs are aggregated at the level of the region, there may be multiple municipality-level locations corresponding to a single firm × occupation × region. If this is the case for a particular job, I calculate z_{ij} as i's *potential commuting distance*: the minimum of the distances to each of these more granular job-locations. Intuitively, if a worker could do the same job at multiple locations within a firm-occupation-region, she would be most likely to work at the location closest to her home.

the cross-sectional matched employer-employee data.

This approach to aggregating the outside option serves two purposes. First, it builds intuition about the content of the relative utility value $\tilde{\mu}_{d(i)jX(i)}$. Intuitively, because logit utility posits that jobs with higher market shares have higher average utilities, jobs that are aggregated into the outside option are very uncommon and thus are likely to be the least desirable alternatives, on average, within the nest. Thus the relative mean utilities identified at the bottom level of the logit structure can be interpreted as utility premia relative to the lowest-quality options. Second, the approach eases computational tractability by effectively truncating the "long tail" of jobs with very few employees. At even the most granular level of the nesting structure, there are hundreds of small jobs that are held by very few workers. This abundance of alternatives makes the maximum likelihood estimation computationally burdensome with minimal benefit – fixed effects for uncommon individual jobs are effectively nuisance parameters that do not substantially affect the second-stage estimation of how wages and amenities contribute to mean utilities. Moreover, because these alternatives are uncommon, their mean utilities will be more noisily estimated, and their characteristics more noisily observed, than more commonly chosen alternatives. Thus, aggregating them should not substantially change the estimation results and provides substantial computational advantages.

Estimation of intermediate- and top-level logits: choices of occupation, and market

After obtaining the scaled relative mean utilities $\tilde{\mu}_{jd(i)X(i)}/\lambda_{mdX}$ and commuting disutilities $\gamma_{mo}/\lambda_{mdX}$, I proceed to the second and third levels of the nested logit structure. For both steps, I first compute the inclusive value of each nest and then estimate level-specific parameters via maximum likelihood.

First, I estimate the parameters of the occupational choice problem conditional on the labor market. The inclusive value of an occupation nest, conditional on the market, is given by

$$\begin{aligned}
I_{imo} &\equiv \lambda_{mdX} \log\left[\sum_{j\in\mathcal{J}_{mo}} \exp[(\mu_{j'd(i)X(i)} + \gamma_{mo}\cdot z_{ij})/\lambda_{mdX}]\right] \\
&= \lambda_{mdX} \log\left[\sum_{j\in\mathcal{J}_{mo}} \exp[(\mu_{j'd(i)X(i)} + \gamma_{mo}\cdot z_{ij} + \mu_{r(mo)d(i)X(i)} - \mu_{r(mo)d(i)X(i)})/\lambda_{mdX})]\right] \\
&= \mu_{r(mo)d(i)X(i)} + \lambda_{mdX} \log\left[\sum_{j\in\mathcal{J}_{mo}} \exp[(\tilde{\mu}_{j'd(i)X(i)} + \gamma_{mo}\cdot z_{ij})/\lambda_{mdX})]\right] \\
&\equiv \mu_{r(mo)d(i)X(i)} + \lambda_{mdX}\tilde{I}_{imo},
\end{aligned}$$

and, correspondingly, the probability that person i chooses occupation o, conditional on being in

market m, is:

$$
\begin{aligned}
s_{io|m} &= \frac{\exp(I_{imo})}{\sum_{o' \in \mathcal{O}_m} \exp(I_{imo'})} \\
&= \frac{\exp[(\mu_{r(mo)d(i)X(i)} + \lambda_{md(i)X(i)} \tilde{I}_{imo})]}{\sum_{o' \in \mathcal{O}_m} \exp[(\mu_{r(mo')d(i)X(i)} + \lambda_{md(i)X(i)} \tilde{I}_{imo'})]}.
\end{aligned} \tag{11}
$$

Intuitively, $\tilde{I}_{imo}$ captures the component of the inclusive value that is identified by the lower level of estimation. These equations clarify that in order to obtain the complete inclusive value, I must also recover the value of the reference alternative in each nest $\mu_{r(mo)d(i)X(i)}$. Note that the mean demographic-specific utilities of the reference jobs in each occupation are recovered from the coefficients on occupation-by-demographic-group fixed effects in this logit.

The last parameter that is recovered from the estimation of the location-level logit is λ_{mdX} for each separately estimated market m, separately for each demographic type d and ability type X. I compute type-specific wages in each by extending the approach of Bonhomme et al. (2019) (henceforth BLM) to estimate worker type proportions and type-specific wages in individual jobs. To do this, I classify jobs into $K = 5$ types and workers into $L = 5$ types within two labor market groups: the Paris high-skill market and all other markets. Jobs are classified using the standard k-means clustering approach, exploiting 20 quantiles of the job's wage distribution. Conditional on the job classification, I estimate worker type proportions within each job type. I report the joint distributions of job and worker types, as well as wage distributions of worker-firm type combinations, in Appendix A.1. Similarly to the results found by BLM, I find evidence of strong sorting of worker types across firm types and minimal productive complementarities. Again, I estimate this level logit using maximum likelihood.

The BLM approach does not identify type-specific wages at individual *jobs*, which I require in order to estimate the skill premia θ_{ol} using the structural wage equation. To identify job-by-type wages, I apply an empirical Bayes algorithm. I use the BLM-estimated job-by-worker-type wage distributions as a prior and compute posterior type distributions using observed wages conditional on job type. Intuitively, knowing the job type's wage distribution allows me to assign type probabilities to individual workers given their observed wages. I can then update the worker type composition within the job and the type-specific wages by integrating over these individual type probabilities. I provide greater detail on the empirical Bayes approach in Appendix A.2. This step yields estimates of worker type proportions and worker-type-specific wages in individual jobs.

Estimation of preference parameters using bottom-level mean utilities

Once I have recovered the μ_{dj} from the sequential estimation of the nested logit model, in the second stage, I estimate the contribution of wages and amenities to these mean utilities. This stage of the algorithm allows me to recover job-specific labor supply elasticities with respect to both wages and amenities. Recall Equation 8, which defines the demographic-specific mean utilities μ_{djX}:

$$\mu_{djX} = \alpha_d \log w_j(X) + \boldsymbol{a}_j' \boldsymbol{\beta_d} + \boldsymbol{y}_j' \boldsymbol{\psi_d} + \xi_{dj}.$$

In the empirical implementation, I focus on estimating differences in preferences parameters for men and women. Then I have ten sets of μ_{djX} corresponding to $d = \{M, W\}$ and $X = 1, ..., 5$.

Recall job j in the data is characterized by a firm k, a market m, an occupation o, and a location l, i.e., $j = (k, m, o, l)$. Both the wage w_j and the observed amenities $\boldsymbol{a}_j$ may be correlated with the unobserved component of job quality ξ_{dj}. To address this source of endogeneity, I require instruments for wages and amenities that are uncorrelated with ξ_{dj}. I construct these instruments by building on the evidence presented in Section 3 that firms' wage amenity policies are correlated across locations and occupations. A candidate approach would instrument for wages w_j with an average of the wages of other jobs at the same firm, $\{w_{j'} : k' = k, (m', o', l') \neq (m, o, l)\}$. However, this leave-out mean approach may not satisfy the exclusion restriction if the firm's unobserved quality shocks are correlated across jobs.

To take the instrument one step further and break this correlation, I construct **predicted values** for both wages and amenities $(\hat{w}_{km'o'l'}, \hat{a}_{km'o'l'})$ for each of the other jobs in firm k. Formally, let $\mathcal{J}_k$ denote the set of jobs at firm k. Each job is indexed by $j = (k, m', o', l')$. For any job $j' \in \mathcal{J}_k$, the predicted value of characteristic $x_{j'} \in \{w_{j'}, a_{j'}\}$ is defined as

$$\hat{x}_{j'} = \mathbb{E}\{x_{j'} | m', o', l'\},$$

estimated via a leave-one-out regression of x on market, occupation, and location fixed effects. Intuitively, these predicted values capture the bundle of wages and amenities that a job would offer to respond to the compensation packages offered by its competitors, not as a result job's unobserved quality shocks. Thus $Z_j^S(x)$ is the average predicted wage or amenity for the firm's other jobs. These instrument strategies mirror Hausman-style "leave-out" instruments in industrial organization, where a product's price in one market is instrumented with its predicted price in other markets.

This design satisfies the two standard assumptions for IV estimation. The instruments are relevant because firms apply broad wage and amenity policies: the predicted values of other jobs in the firm are strongly correlated with the focal job's characteristics. The instruments satisfy the

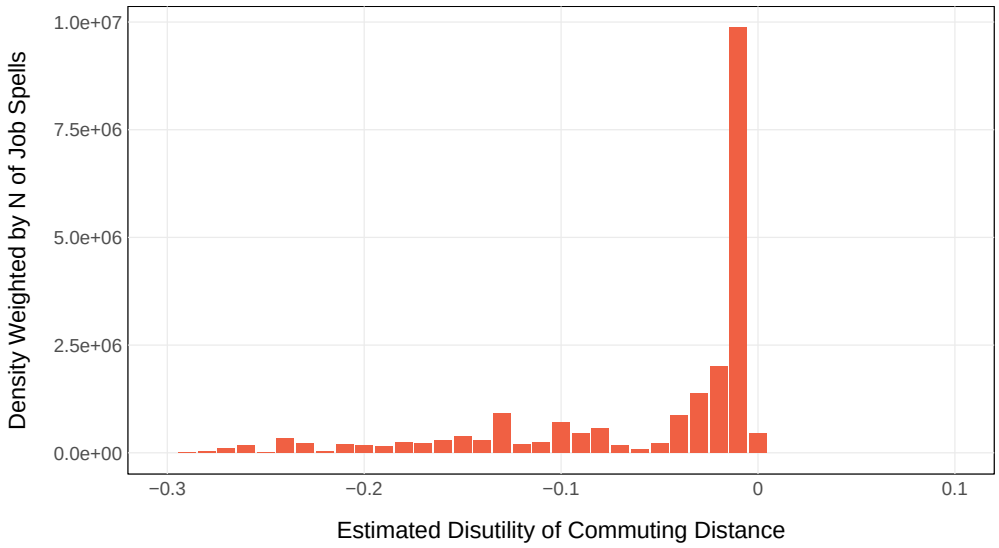
exclusion restriction because by predicting each other job's characteristics using only external market, occupation, and location fixed effects, the instruments purge firm- or job-specific unobserved quality that could be correlated across jobs within the firm. Averaging these predictions then produces a shifter that reflects the outside option environment of all the jobs at the firm, which should be uncorrelated with the unobserved quality of the focal job.

4.2. Results

Disutility of commuting distance

Figure 6 plots the distribution of estimates of γ_{mo} for each granular nest (m, o), obtained by multiplying the estimated coefficient by the nesting parameter λ_m. The vast majority of the estimates are negative, as expected: longer commutes should make jobs less attractive, all else equal. Out of 278 market-occupation nests, I estimate a positive distance coefficient in only one case. This is comparable to the result of Azar et al. (2022), who find one such positive and statistically significant estimate among 822 markets. Moreover, the single positive estimate is equal to 0.001 after scaling by the substitution parameter and is statistically indistinguishable from zero. I weight the density by the number of workers in each nest, so that the density shows the number of workers that are associated with a disutility of commuting in each bin rather than the number of nests with a given estimate. Thus, the figure summarizes the distribution of worker preferences across the economy.

FIGURE 6. Estimates of Disutility of Commuting Distance by Granular Nest

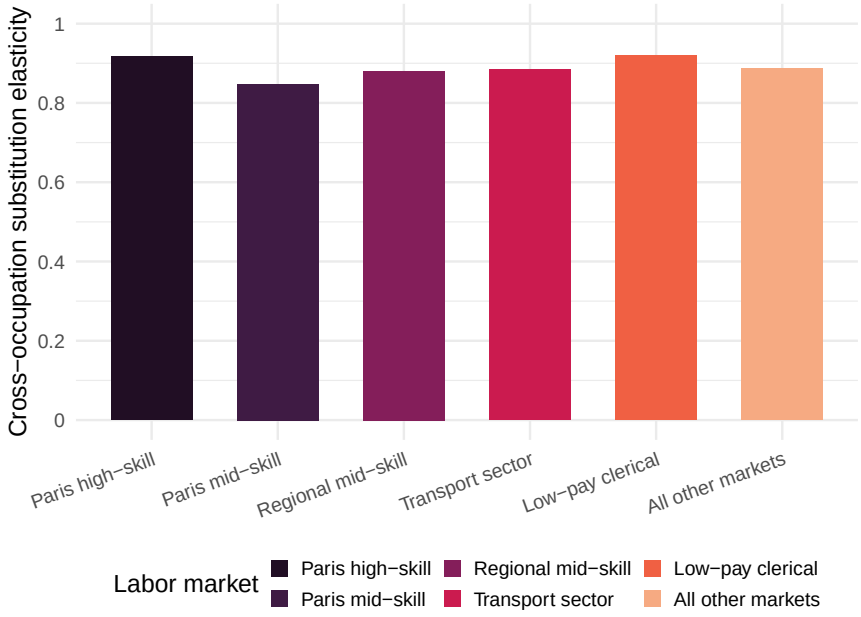


Notes: This figure plots a histogram of the estimated parameters γ_{mo} for each granular labor market × 2-digit occupation nest. Estimates are derived from the coefficient γ_{mo}/λ_m at the bottom level of the nested logit structure as shown in Equation 9 and adjusted by the market-specific elasticity of substitution across occupations λ_m obtained at the upper level of the nested logit structure. The histogram is weighted by the number of job spells in each nest.

Substitution between occupations within a labor market

Figure 7 summarizes the estimates of the market-specific elasticity of substitution λ_m. I estimate λ_m for each of the twenty data-driven labor markets. To summarize the estimates for markets with similar compositions, the "all other markets" bin shows the weighted average estimate across the remaining fifteen markets, where markets are weighted by the number of workers they represent.

FIGURE 7. Estimates of Occupational Substitution Elasticities λ_{mo} by Market



Notes: This figure reports estimates of the occupational substitution parameter λ_{mo} corresponding to each labor market. Estimates are obtained from the coefficient on the inclusive value from GLM estimation of the top level of the nested logit structure described in Equation 11. The figure plots estimates for five individual markets that differ substantially from the others in terms of their earnings distributions and occupational/sectoral compositions. The remaining fifteen markets are aggregated into the "all other markets" bin, which plots the average within-market substitution elasticity across these markets. In this bin, market estimates are weighted by the number of workers in the market to construct the average.

The average estimates of λ_m for each market fall between zero and one, a restriction implied by the nested logit model but not guaranteed by the estimation method. Thus the estimation results do not reject the nested logit model structure. All of the estimates reject the null of $\lambda_m = 1$, but most are economically close to it: 89 percent of workers have a market-gender-type-specific λ_{mgX} between 0.8 and 1.0. These estimates indicate that workers do not view jobs in different occupations as highly differentiated within their choice sets. Formally, knowing the 2-digit occupation of a job is somewhat informative about a worker's preference for it, but not strongly so.

These estimates are much larger than the market substitution parameters estimated by Azar et al. (2022). They estimate elasticities of substitution between commuting zone × 6-digit occupation

markets, with estimates in the range of 0.097 to 0.159. Three factors account for the difference. First, I condition workers' choice sets on both unobserved heterogeneity and on markets defined at a coarser level of aggregation. Workers rarely move across markets as I define them, consistent with a substitution elasticity close to zero at the market level. Second, conditional on unobserved type, workers cluster in a relatively narrow set of occupations. For example, the highest-type workers in the Paris high-skill market hold highly-skilled white-collar jobs 96% of the time. Given this concentration, occupation provides little additional information about how substitutable two jobs are once the choice set is determined. Third, and relatedly, the finding that occupations are not highly differentiated within choice sets suggests that the partitioning successfully captures how workers perceive different alternatives as inside or outside their consideration sets, regardless of the specific occupation assigned to a given job.

Wage and amenity elasticities for men and women

Table 5 reports the estimates of the coefficients α and β on wages and amenities in Equation 8 using as instruments the predicted wages and amenities for other jobs at the same firm. In each case, predicted wage for a given job is defined as the leave-out average wage of other jobs in that market and occupation, leaving out the own job. Predicted amenities are defined analogously. The F statistics illustrate that across all specifications, the instruments have extremely strong first stages. The columns of the table use different combinations of controls and fixed effects: columns (1) and (2) use only market × year fixed effects; columns (3)-(4) use market × year and sector × year fixed effects; and (5)-(6) adds controls for a second-order polynomial in firm size. The firm size controls aim to account for the fact that some jobs have a large market share not because they are strongly preferred by workers but because they hire large numbers of workers. I take the market-sector specification (columns 3–4) as preferred. This specification includes the richest set of controls with significant estimates for men's and women's preferences that still preserves variation across firms, which better represents how workers evaluate and choose among jobs in practice.

Across all specifications, the coefficient on log earnings is positive and statistically significant, confirming that wages are an important determinant of job utility. A consistent gender gap emerges in wage sensitivity: the earnings coefficient ranges from 1.350 to 1.617 for men but only 0.704 to 1.110 for women. Men's utility is more responsive to earnings than women's across every identification strategy.

The coefficients on schedule variability display a different pattern. Men appear to have a positive willingness to pay for schedule variability in the specifications without firm fixed effects, with coefficients ranging from 0.625 to 0.682. However, this effect disappears entirely once firm fixed effects are included, with the coefficient falling to 0.126 and losing statistical significance. This

TABLE 5. Impact of wages and amenities on utility derived from jobs: instruments based on average predicted wages and amenities of other jobs at the same firm

Gender	**Men**	**Women**	**Men**	**Women**	**Men**	**Women**	**Men**	**Women**
Mean log earnings	1.478***	0.704***	1.612***	0.951***	1.350***	0.905***	1.617***	1.110***
	(0.327)	(0.202)	(0.335)	(0.177)	(0.307)	(0.168)	(0.276)	(0.201)
Schedule variability	0.682***	-0.105	0.678***	-0.050	0.625***	-0.119	0.126	0.008
	(0.165)	(0.095)	(0.148)	(0.135)	(0.112)	(0.119)	(0.109)	(0.115)
Greedy hours	-0.263	-0.417***	-0.566**	-0.435***	-0.241	-0.327**	-0.261**	-0.161***
	(0.268)	(0.117)	(0.253)	(0.151)	(0.224)	(0.150)	(0.110)	(0.060)
Market × year FE	X	X	X	X	X	X		
Sector × year FE			X	X	X	X		
Firm size controls					X	X		
Firm FE							X	X
F-stats								
Mean log earnings	270576	163244	226410	130497	226621	130486	180478	96855
Schedule variability	3906	2170	1831	1168	1724	1150	533	337
Greedy hours	5482	3989	3996	2943	3856	2906	2893	1980
N	100575	80616	100575	80616	100575	80616	100575	80616

Notes: This table presents estimates of 2SLS regressions of the mean utility of each job on the job's mean log earnings as well as its mean "schedule variability" and "greedy hours" measures as defined in section 2. Wages and amenities are instrumented with the predicted wages for other jobs at the same firm. The predicted wage and amenity for job j' at the same firm is computed as the mean value of those characteristics for jobs $\tilde{j}$ at other firms in the same market, occupation, and location, excluding job j'. Regressions are estimated separately for men's and women's mean utilities. Standard errors are clustered at the firm level, allowing for correlations in unobserved quality across jobs within the firm.

attenuation suggests that the positive association between schedule variability and utility in the baseline specifications reflects correlation with unobserved firm-level characteristics rather than a genuine preference for variable schedules. Women's coefficients on schedule variability are small and statistically indistinguishable from zero across all specifications.

The estimates for preferences over "greedy hours" reveal a more nuanced pattern. While the point estimates of the disutility of greedy hours are similar for men and women in absolute terms, and in some specifications women's coefficient is larger while in others it is smaller, the implied willingness to pay (WTP) to avoid greedy hours is consistently higher for women in the three main specifications. Expressed as a share of earnings, women's WTP to reduce greedy hours by one standard deviation ranges from 0.36 to 0.59, compared to 0.18 to 0.35 for men. This difference arises primarily from women's lower sensitivity to wages rather than from a uniformly stronger distaste for demanding schedules. The finding is consistent with prior evidence that women exhibit a greater willingness to trade earnings for amenities, while highlighting that the underlying channel may be differences in the marginal utility of income rather than stronger intrinsic preferences for non-wage job characteristics *per se*. The one exception to this pattern appears in the firm fixed effects specification, where men's and women's WTP estimates converge (0.16 and 0.15, respectively). I interpret the firm fixed effects

results primarily as a robustness check on the instrument validity rather than as a preferred estimate of worker preferences, since this specification absorbs the across-firm variation that is central to how workers choose among jobs.

Having estimated the gender-specific coefficients and the substitution parameters from the nested logit estimation procedure, I can now derive job-specific labor supply elasticities for men and women. I am primarily interested in the *elasticity of labor supply within market m to job j*. Intuitively, this elasticity captures how labor supply to the job changes for workers who are already within the market. Because cross-market substitution elasticities are very small and statistically indistinguishable from zero, I assume that firms do not consider how changes in wage and amenity policies will induce workers to switch markets entirely. This elasticity of labor supply to job j of workers in demographic group $d = \{M, F\}$ with respect to the wage at the same job is given by:

$$\varepsilon^{w}_{dj} = \frac{\alpha_d}{s^{d}_{j|m} N_{dm}} \sum_{i=1}^{N_{dm}} s_{ij|m} \left[\frac{1}{\lambda_m} + \frac{\lambda_m - 1}{\lambda_m} s_{ij|mo} - s_{ij|m} \right], \tag{12}$$

where α_d is the wage coefficient from Table 5; N_{dm} is the number of workers of demographic type d in market m; $s^{d}_{j|m}$ is the observed market share of job j in market m for workers in demographic d; λ_m is the elasticity of substitution between occupations for market m; $s_{ij|m}$ is worker i's predicted probability of choosing job j in occupation o conditional on market m; $s_{ij|mo}$ is i's predicted probability of choosing job j conditional on market m and occupation o. I provide a full derivation of the elasticity formula in Appendix C.3.

TABLE 6. Average wage and amenity elasticities by market and gender

	Wage Elasticities		**Amenity Elasticities**	
Market	**Men**	**Women**	**Men**	**Women**
Paris high-skill	1.756	1.065	0.617	0.487
Paris mid-skill	1.960	1.126	0.688	0.515
Regional mid-skill	1.830	1.101	0.642	0.504
Low-pay clerical	1.787	1.032	0.627	0.472
Transport sector	1.841	1.185	0.646	0.542
All other markets	1.888	1.089	0.663	0.498

Notes: This table shows quantiles of the estimated distributions of labor supply elasticities with respect to the wage and the flexibility amenity across markets. Elasticities are measured at the level of the individual job for each gender. Averages are weighted by the number of workers in each job.

Table 6 summarizes the average job-level elasticity estimates by gender in each of the major

markets. Columns (1)–(2) report wage elasticities and columns (3)–(4) report elasticities with respect to the greedy hours disamenity, with the sign normalized to reflect a positive labor supply response to an improvement in the amenity.

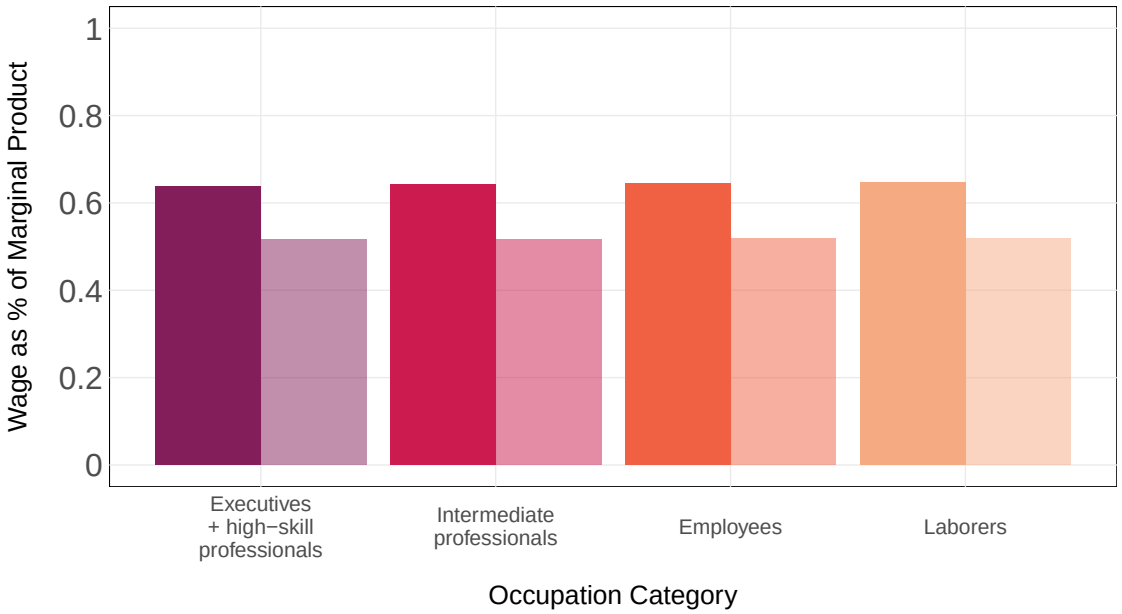
The table highlights two key patterns. First, wage elasticities are systematically lower for women than for men across all markets, implying that women face higher wage markdowns. Men's wage elasticities range from 1.76 to 1.96 across markets, implying markdowns of 36–39 percent; women's wage elasticities range from 1.03 to 1.19, implying markdowns of 48–51 percent. The gender gap in markdowns is thus on the order of 10–15 percentage points, driven primarily by women's lower sensitivity to wages in utility terms rather than by differences in the competitiveness of the markets in which men and women work. Indeed, the elasticity estimates are reasonably uniform across markets, with no clear gradient by skill level. The Paris high-skill and low-pay clerical markets exhibit the lowest elasticities and highest markdowns, reflecting larger estimates of the nesting parameter λ_m in these markets, potentially driven by greater stratification across worker types.

Second, amenity elasticities follow a similar pattern: men's elasticities with respect to greedy hours range from 0.62 to 0.69, while women's range from 0.47 to 0.54. Although women have a higher willingness to pay to avoid greedy hours (as shown in Table 5), men's amenity elasticities are actually higher in absolute terms because men's greater wage sensitivity amplifies their labor supply response to amenity differences.

These elasticity estimates are meaningfully smaller than those reported by Azar et al. (2022), and consequently imply larger markdowns. Our coefficient estimates are similar to theirs; the difference is driven by the nesting parameters λ_m and by differences in market definitions and substitution structure. The gap may also reflect a distinction between application and employment elasticities: Azar et al. (2022) estimate application elasticities and extrapolate to labor supply, but workers conducting a job search may exhibit more correlated preferences in their applications than in their actual job choices conditional on entering a specific market. A direct comparison would require accounting for unobserved worker heterogeneity in their framework.

Figure 8 plots the average markdowns implied by the wage elasticity estimates for men and women in each of the four occupational skill categories. The markdown estimates highlight two results. First, because elasticities are similar across occupational skill, markdowns are also uniform on average conditional on gender: about 36 percent for men and about 48 percent for women in each occupation. Second, women face larger markdowns than men in each occupation because they have lower job-specific labor supply elasticities. These markdowns reflect consistent differences across jobs rather than compositional differences in the types of jobs held by men and women.

FIGURE 8. Average Markdowns By Occupation and Gender



Notes: This figure plots wage markdowns across 1-digit occupation groups. Markdowns are calculated for each job $j \times$ gender $g \times$ worker type X cell as $\varepsilon^w_{jgX}/(1+\varepsilon^w_{jgX})$, where ε^w_{jgX} is (g,X)'s elasticity of labor supply with respect to the wage at job j. Men's markdowns are plotted in dark colors (left side of each group) and women's markdowns are plotted in light colors (right side of each group). Occupation-level averages are computed by weighting job-gender-type-level estimates by the number of workers in the job-gender-type. The sample is the subset of the BTS-Postes 2017-2023 (excluding 2020) that is matched to the Labor Force Survey (EEC) at the job $\times$ year level.

5. Labor demand: how do firms set wages and amenities?

In this section, I develop and estimate a model of labor demand that builds on the stylized model presented in Section 3.1. I then use the model to recover productivity penalties associated with providing amenities and input elasticities of substitution between labor from different occupations. I estimate the model via two-step GMM using microdata on firms' choices of employment, wages, and amenities in each job conditional on the wage and amenity elasticities estimated in Section 4. After obtaining the parameter estimates, I decompose sources of variation in the structural wage and amenity equations to quantify the contributions of technology and preference channels to firms' amenity choices.

5.1. Demand for occupation-specific labor with imperfect substitution

The purpose of the labor demand model is to estimate the parameters governing firms' technology of amenity provision and the complementarity or substitutability of occupation-specific labor inputs. As I establish in the stylized model of Section 3.1, these parameters are key inputs into the technology

and preference forces that drive the relationship between wages and amenities. The signs of the covariances between these different components sheds light on how we can expect wages and amenities to be positively or negatively correlated.

The model has three key features that generalize the stylized framework. First, the CES labor technology allows for imperfect substitution between a variety of occupations. Imperfect substitution generates interdependence in the firm's choices of wages and amenities and is important for estimating firms' wage-amenity cost tradeoff. Second, firms choose a vector of amenities that may incur different productivity penalties in different occupations. This allows for heterogeneity in the productivity cost of providing different types of amenities. Third, I allow for unobserved heterogeneity in worker types. I establish identification of occupation-specific wage premia for discrete worker types conditional on estimated type-specific wages in each job. After identifying skill premia, I can explicitly control for unobserved heterogeneity in the estimation of the remaining labor demand parameters.

Workers have a finite distribution of skill types $l = 1, .., L$. Within an occupation o, each type commands a *skill premium* θ_{lo} that parameterizes the effective units of labor supplied by a worker of skill type l in occupation o. Then the labor supply to occupation o at firm j in time t is a sum over types, weighted by skill premium:

$$L_{ojt} = \sum_{l=1}^{L} \theta_{ol} L_{lojt}.$$

I assume that the skill premium θ_{lo} is common across firms and abstract away from differential returns to skill by firm.

At time t, the firm chooses wages w_{lojt} for each occupation × worker type and a vector of amenities $\boldsymbol{a_{ojt}}$ for each occupation $o \in \mathcal{O}_j$. The firm can choose different wages for each skill type, but must provide the same amenity to all workers in occupation o. The production function has returns to scale in labor α_j, where labor is a CES aggregate over occupation-specific inputs:

$$\begin{aligned} Y_{jt} &= A_{jt} \cdot L_{jt}^{\alpha_j}, \qquad \text{where} \\ L_{jt} &= \left(\sum_{o \in \mathcal{O}_j} (z_{ojt} L_{ojt})^{\rho_j} \right)^{\frac{1}{\rho_j}}, \\ z_{ojt} &= \exp(\boldsymbol{\delta}'_{\boldsymbol{oj}} \boldsymbol{a_{ojt}}), \qquad \text{and} \\ L_{ojt} &= \sum_{l=1}^{L} \theta_{ol} L_{lojt}. \end{aligned} \tag{13}$$

As before, δ_{oj}^{k} represents firm j's productivity penalty associated with providing amenity k to workers in occupation o, and ρ_j determines the elasticity of substitution between occupations. I assume that there is no direct cost of providing the amenities $\boldsymbol{a}_{ojt}$.[11]

The firm maximizes profit by choice of wages and amenities, subject to the constraint that labor supply equals labor demand:

$$\max_{\{w_{lojt},\boldsymbol{a_{ojt}}\}_{l,o}} \left\{ A_{jt} \cdot \left(\sum_{o \in \mathcal{O}_j} (z_{ojt} L_{ojt})^{\rho_j} \right)^{\frac{\alpha_{Lj}}{\rho_j}} - \sum_{o \in \mathcal{O}_j} \sum_{l \in \mathcal{L}_{oj}} w_{lojt} L_{lojt} \right\},$$
$$\text{s.t. } L_{lojt} = S_{lojt}(w_{lojt}, \boldsymbol{a}_{ojt}).$$

The first order conditions for each of the type-by-occupation-specific wages w_{lojt} yield the log wage equations:

$$\log(w_{lojt}^{*}) = \log\left(\frac{\varepsilon_{ojt}^{w}}{1+\varepsilon_{ojt}^{w}}\right) + \psi_{jt} + (\rho_j - 1)\log L_{ojt} + \rho_j \boldsymbol{\delta'_{oj}} \boldsymbol{a_{ojt}} + \log\theta_{ol} + \epsilon_{ojt}^{w}, \quad \text{where}$$
$$\psi_{jt} = \log\alpha_{Lj} + \log\left(\frac{Y_{jt}}{L_{jt}}\right). \tag{14}$$

Note that the term $\log\left(\frac{Y_{jt}}{L_{jt}}\right)$ is constant across all observations at firm j and time t and the log labor share $\log\alpha_{Lj}$ is assumed to be constant at the firm level. These two terms can thus be combined and expressed as a firm-period fixed effect ψ_{jt}. The firm-period fixed effect captures any productivity shocks that are common to all workers at firm j at time t, so that I do not need to make any distributional assumptions on the productivity shocks to eliminate bias from this source. Furthermore, ε_{ojt}^{w} has already been estimated from the labor supply side of the model, so I treat markdowns as known and estimate demand parameters conditional on them.

As before, the firm chooses amenities to attract more labor at a lower wage, taking into account that the amenities may carry penalties to the productivity of the labor inputs. To choose the optimal

[11] This assumption is reasonable in the context of the limited set of amenities related to scheduling and demands on workers' time on which this paper focuses. Direct costs related to scheduling may involve some setup cost (e.g., software to facilitate shift coordination or a small home office subsidy), but these are typically fixed investments with minimal marginal usage cost per worker. Thus, these costs would not enter the static optimization problem after fixed costs have been paid. However, the assumption may not generalize to other amenities that are not considered in this paper, such as healthcare subsidies with a clear and non-trivial marginal cost per worker.

level of each amenity k, the firm equates marginal returns and marginal labor supply costs:

$$\log a_{ojt}^{k*} = \log\left(\frac{\varepsilon_{ojt}^{a^k}}{1+\varepsilon_{ojt}^{w}}\right) - \log(-\delta_{oj}^{k}), \tag{15}$$

where $\varepsilon_{ojt}^{a^k}$ is the elasticity of labor supply to job oj with respect to amenity k. The equilibrium condition for the amenity is a generalized version of the same condition in the case with a single amenity. Full derivations of the equilibrium conditions for wages and amenities are provided in Appendix C.4.2.

To estimate the structural equations, I specify mean-zero error terms ϵ_{ojt}^{w} and ϵ_{ojt}^{a} that capture idiosyncratic deviations from the optimal wage- and amenity-setting rules conditional on the observed wages, elasticities, firm-year fixed effects, labor supply, and amenities specified in Equations 14 and 15. I estimate the model using cluster-robust standard errors, clustered at the market-occupation level. This allows for arbitrary correlations in unobserved shocks within the firm while maintaining consistency of the estimator under minimal assumptions about the error structure. The only required assumption is that for some instruments Z_{ojt}^{w} and Z_{ojt}^{a}:

$$\mathbb{E}[\epsilon_{ojt}^{w}|Z_{ojt}^{w}] = 0,$$
$$\mathbb{E}[\epsilon_{ojt}^{a}|Z_{ojt}^{a}] = 0.$$

Table 7 summarizes the parameters estimated on the demand side of the model and the sources of variation that identify them. The parameters are estimated at different levels of granularity. The labor share of output α_{Lj} is estimated for 18 different sectors, allowing for heterogeneity in the role of labor in production across different types of economic activity. The substitution parameter ρ_j is estimated at the level of the sector × job class $k = 1, ..., 5$. The skill premia are estimated separately by gender at the level of the 2-digit occupation × worker type $l = 1, ..., 5$, and are identified up to a normalization by the skill premium of the lowest worker skill type within each occupation. Finally, amenity costs are estimated by 1-digit occupation and job class k. Next, I discuss how the parameters are identified and specify key choices that allow me to bring the model to data, beginning with the definition of worker and job types and the estimation of skill premia for each.

5.2. Estimation of skill premia θ_{ol} for discrete unobserved worker types

I estimate the skill premium parameters θ_{ol} associated with specific worker types and occupations in two steps. In the first step, I estimate worker type proportions and worker-type-specific wages in individual jobs using the approach of Bonhomme et al. (2019) as described in Section 4.1. In

TABLE 7. Summary of Labor Demand Parameters and Sources of Identification

Name		Unique Parameters	Mean Estimate	Moments Used for Identification
Labor share	α_{Lj}	18	0.669	Wage equation—sector FE
Substitution parameter	ρ_j	65	0.986	Wage equation–Bartik instruments
Skill premium (men)	$\theta_{ol,m}$	120	1.93	Wage equation—occ. × worker type FE
Skill premium (women)	$\theta_{ol,f}$	120	1.73	Wage equation—occ. × worker type FE
Amenity cost	δ_{oj}	40	–0.389	Amenity equation—sector × firm type FE

Notes: This table summarizes the parameters of the labor demand model and the data moments used to identify them. Skill premia θ_{ol} are identified up to a normalization so that the lowest skill category has a skill premium of 1 (log skill premium of 0).

the second step, after identifying job-by-type wages, I estimate type premia using a fixed effect regression. Because I can treat the wage markdown term as known, I can express the structural wage equation in terms of fixed effects. This allows me to pre-estimate θ_{ol} parameters using fixed effects, reducing the number of parameters that need to be estimated using GMM. The wage equation can be rewritten:

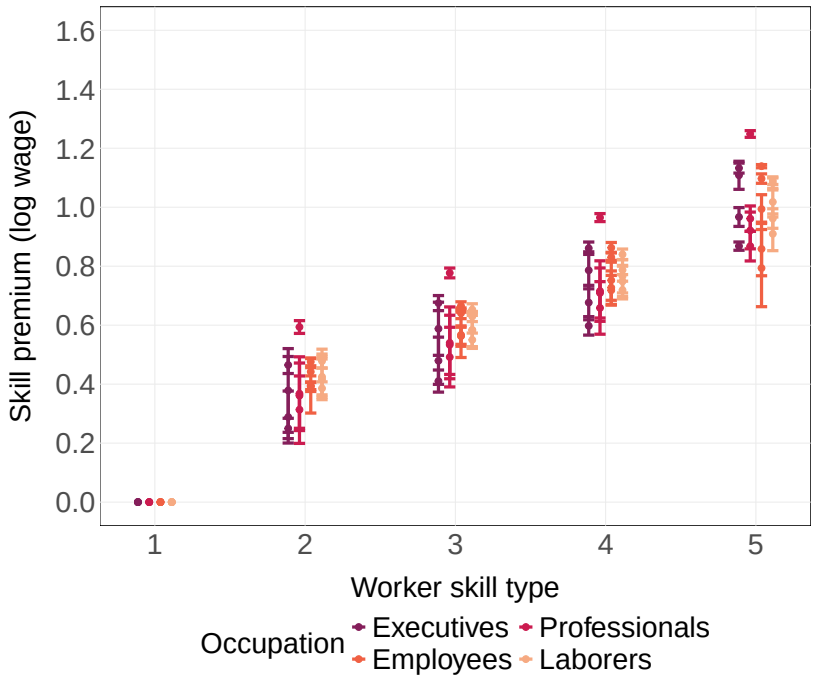
$$\log(\tilde{w}^*_{lojt}) = \log \alpha_{Lj} + \psi_{jt} + \gamma_{ojt} + \log \theta_{ol} + \epsilon^w_{ojt}, \tag{16}$$

where $\log(\tilde{w}^*_{lojt}) \equiv \log(w^*_{lojt}) - \log\left(\frac{\varepsilon^w_{ojt}}{1+\varepsilon^w_{ojt}}\right)$ is the residualized log wage, $\psi_{jt} = \log(Y_{jt}L_{jt}^{-1})$, and $\gamma_{ojt} = (\rho_j - 1)\log L_{ojt} + \rho_j \boldsymbol{\delta}'_{\boldsymbol{oj}} \boldsymbol{a}_{\boldsymbol{ojt}}$. This equation clarifies that $\log \theta_{ol}$ is identified from a type-by-occupation fixed effect. Skill premia in each occupation are identified relative to a reference skill group, which I define to be the lowest skill category. I can absorbe the job-year fixed effect γ_{ojt} because this component is common to all worker types within the same (o, j, t) cell. Because there are multiple worker types within jobs, I can identify the fixed effect parameters θ_{ol} even with these granular controls.

Figure 9 reports estimates of log skill premia by worker skill type $k = 1, ..., 5$ and 2-digit occupation. Estimates are color-coded by 1-digit occupation, corresponding to the skill category of the profession.[12] The figure shows that each worker skill type—ordered from lowest-paid to highest-paid— commands roughly consistent pay premia across a wide variety of occupations. In most occupations, the second worker category is paid 0.3-0.5 log points more than the lowest category; the third category is paid a premium of 0.5-0.7; the fourth category earns a premium of

[12]The vast majority of occupations are classified into four groups: "executives and higher intellectual professions", "intermediate professionals", "employees" and "laborers". I omit the small share of workers classified as farmers, artisans, and traders.

FIGURE 9. Estimates of Skill Premia θ_{ol} by Occupation and Worker Skill



Notes: This figure plots means and standard errors of parameter estimates of skill premia θ_{ol} obtained from the fixed effect regression specified in Equation 16. Parameters are estimated by worker skill type $l = 1, ..., 5 \times$ 2-digit occupation, where occupation is measured using the French PCS-ESE classification. Occupations are color-coded using the 1-digit occupation under the same classification. The sample is the subset of the BTS-Postes 2017-2023 (excluding 2020) that is matched to the Labor Force Survey (EEC) at the job level. Standard errors are clustered at the firm level.

0.7-0.9, and the highest category is paid a premium ranging from 0.8-1.2 log points. The estimates imply that the highest-quality workers are up to three times as productive as the lowest-quality workers. In Appendix A.3, I show estimates of skill premia separately by gender. Across and occupations, women earn consistently lower pay premia than men in the same skill category. This skill premium could reflect true differences in productivity, but it could also encode systematic pay differences that women experience throughout their careers for other reasons, such as discrimination.

5.3. Estimation and identification of amenity costs δ_{oj} and substitution parameter ρ_j

To bring the model to data, I need to specify the level at which the ρ_j and δ_{oj} parameters vary. In a setting with infinite data and computing power, the ideal estimator would recover firm-specific parameters. However, we must account for computing and precision constraints while maintaining sufficient heterogeneity to learn about the relationships between these parameters and other covariates of interest, such as sector, skill, and firm productivity.

To balance these competing objectives, I discretize firm heterogeneity to recover rich patterns

in firm technologies that inform the strength of compensating versus agumenting differentials. Rather than estimate firm-specific technology parameters, I group jobs together by sector s and 1-digit occupation o. In the French PCS-ESE occupation classification scheme, the first digit of the occupation code corresponds to the skill level of the job. Thus, this level of granularity allows us to learn about how the productivity costs of amenity provision vary with the skill of the worker. Additionally, I use the job classifications developed in the identification of the θ_{ol} parameters to allow for the possibility that high-paying jobs have different substitution patterns even within the sector.

With 13 sectors, 4 occupational skill groups, and 5 job types, I estimate a total of 65 ρ_{sk} substitution parameters and 40 δ_{ok} parameters. Because I take advantage of rich microdata at the job level, discretizing firm heterogeneity strengthens model identification by ensuring that the number of moment conditions far exceeds the number of parameters to be estimated. I focus on estimating parameters for just one amenity: the greedy-hours composite for which workers have a clear distaste in the supply-side estimation results.

After discretizing job heterogeneity, the amenity equation that I bring to data becomes:

$$\log a^*_{ojt} = \log\left(\frac{\varepsilon^a_{oj}}{1+\varepsilon^w_{oj}}\right) - \log(-\delta_{os}) + \epsilon^a_{ojt},$$

where ϵ^a_{ojt} is a mean-zero error term as specified above. Because ε^w_{oj} and ε^a_{oj} have already been estimated on the supply side of the model, the only parameters remaining to be identified from this equation are the δ_{os}. The structural equation implies that I can identify these parameters as coefficients on group fixed effects. Importantly, this implies that I do not require instruments for amenities to identify these parameters.

To identify the substitution parameters ρ_s, I require instruments for endogenous labor inputs L_{ojt} in the wage equation. I instrument for employment in job $j = (k, m, o, s, t)$—where k is the firm, m is the market, o is the occupation, and s is the sector—with growth in the same occupation-sector (o, s) outside of the focal market. For each market-occupation-sector (m, o, s) at time t, I compute

$$Z_{most} \equiv \Delta \log(L_{ost} - L_{most}),$$

i.e., the year-over-year growth in the log of employment in (o, s) outside of market m. The instrument captures demand shocks for the job's occupation-sector outside the focal market. Because this instrument is constant at the market-occupation level, I sacrifice some cross-firm variation. However, since I only need to use the instruments to identify sector-wide parameters ρ_s, I still have more variation than is strictly needed to identify all of the parameters. The identifying assumption is that,

conditional on the firm-year fixed effects ψ_{jt}, the instrument affects own-job wages only through its impact on L_{ojt} and is orthongonal to job-specific wage shocks that are distinct from the firm-year productivity effects. The moment conditions for GMM estimation are then:

$$\mathbb{E}\left[r^w_{ojt}(\rho^c, \delta^c_{ok}) \cdot Z^w_{ojt}\right] = 0 \quad \text{and}$$
$$\mathbb{E}\left[r^a_{ojt}(\rho^c, \delta^c_{ok}) \cdot Z^a_{ojt}\right] = 0,$$

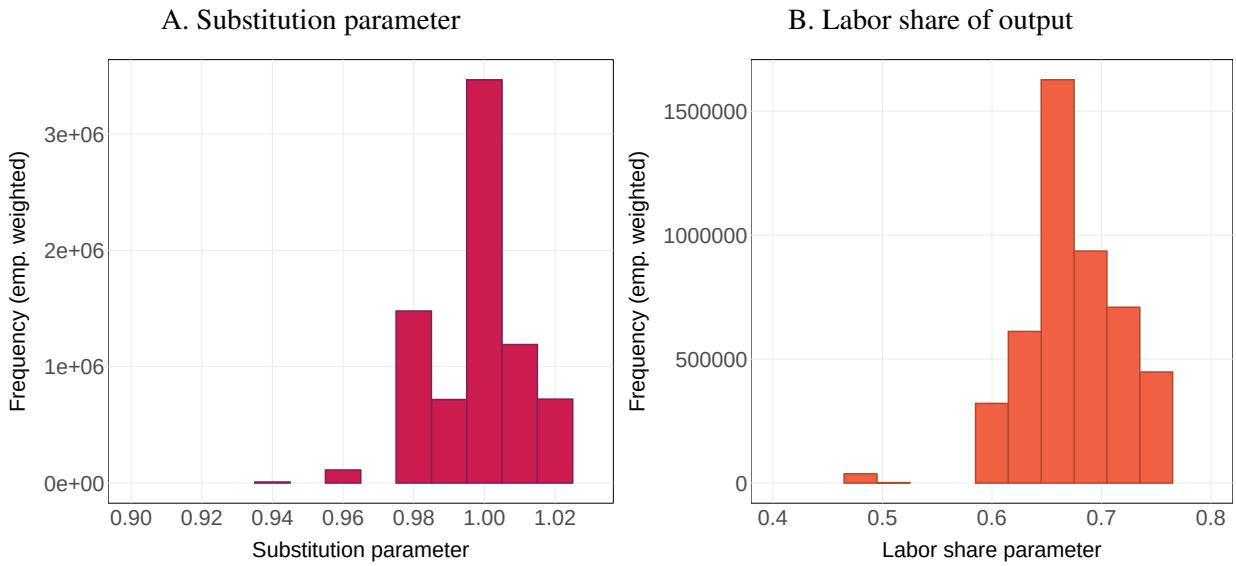
where Z^w_{ojt} denotes the excluded instruments for wages and Z^a_{ojt} is a group (s, o) indicator. I estimate the system via two-step GMM to compute optimal weights and obtain standard errors using 200 iterations of a cluster-robust block bootstrap.

5.4. Estimation Results

Figure 10 reports estimates of the cross-occupation substitution parameters ρ_{sk} and labor shares of output α_s. Substitution parameters are estimated at the level of the sector s × BLM firm class $k = 1, ..., 5$, and labor shares are estimated by sector. Panel A shows that all of the ρ_{sk} estimates are close to one: the largest point estimate, for fourth-category firms in the information and communications sector, is just 1.03; the smallest is 0.92, for second-category firms in the "other services" sector. While many estimates are significantly different from the Cobb-Douglas benchmark, the magnitude of the difference is economically small. These results suggest that firm-wide wage and amenity setting is unlikely to be driven by strong technological complementarities across occupations. Panel B shows that labor shares range from 0.49 for the extractive sector to 0.76 for "other services." While these parameters are not the primary focus of the estimation, they are consistent with other estimates of sector-specific labor shares from the literature.

Figure 11 shows the absolute values of the estimated amenity cost $\exp(\delta_{ok})$ to provide the flexibility amenity, measured as the penalty to the worker's productivity associated with a 1-standard-deviation increase in the flexibility amenity. Parameters are estimated separately for the Paris high-skill markets and all other markets at the level of the firm type $k = 1, ..., 5$ × 1-digit occupation o—a close proxy of worker skill. Occupations are differentiated by color, so that the highest-skill occupation group ("executives and high-skill professionals") is in dark purple and the lowest-skill group ("laborers") is in pale orange. The figure shows that there is a very high cost to limiting greedy hours (or, conversely, a productivity benefit to allowing greedy hours) in the highest-skill occupations. A one-standard-deviation increase in the level of the amenity results in workers being 50-73 percent as productive. The magnitude of the penalty δ_{ok} shrinks as skill level decreases: the analogous estimates for intermediate professionals are 76-81 percent, for employees are 75-85 percent, and for laborers are 86-88 percent. Higher-skill workers thus face steeper penalties to their

FIGURE 10. Histograms of Estimated Production Parameters

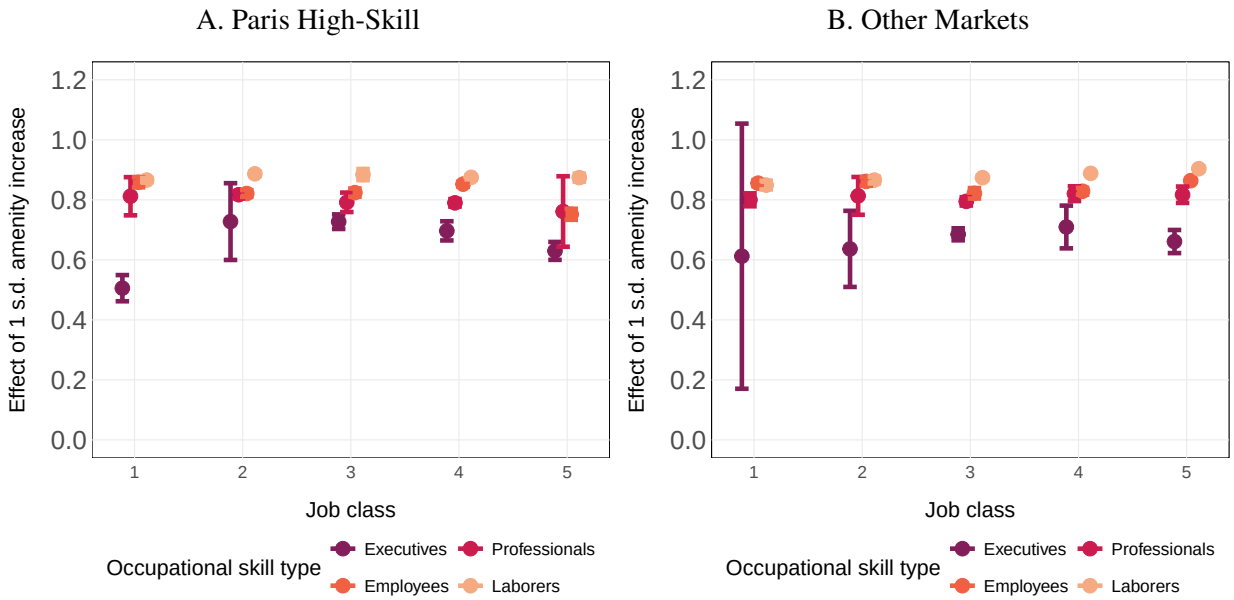


Notes: This figure plots the distributions of parameter estimates of elasticities of substitution across occupations ρ_j (Panel A) and distributions of labor share of ouput α_j (Panel B). Substitution parameters are estimated by job class $k = 1, ..., 5 \times$ sector, where sector is measured as the 21-digit principal activity code. Labor shares are estimated for each sector. The sample is the subset of the BTS-Postes 2017-2023 (excluding 2020) that is matched to the Labor Force Survey (EEC) at the job level.

productivity than lower-skill workers do. Within skill group, there is relatively little variation in the penalty, with the exception of the highest-skill workers. Among the highest-skill group, workers in the lowest-paying jobs have the steepest costs of flexibility, followed by workers in the highest-paid jobs.

These estimates highlight that differences in amenity provision are not simply a matter of firms weighing wages against non-wage benefits. Instead, the technological cost of providing certain amenities—here, limiting work outside regular hours—varies systematically across occupations. Variation across markets is a direct consequence of these differences by skill. High-skill and high-pay jobs exhibit much larger penalties to offering predictable or limited schedules, suggesting that productivity in these settings is more tightly linked to workers' availability beyond standard hours. In other words, working in competitive and productive jobs comes with a greater technological cost of providing desirable hours. This pattern suggests that compensating differentials for hours and demands are not uniform across the economy, but instead reflect underlying technological differences in how productivity is generated in different labor markets.

FIGURE 11. Estimated productivity cost of limiting greedy hours by skill level and sector



Notes: This figure plots estimates of the amenity cost parameters δ_{oj}, estimated separately for the Paris high-skill market and all other markets pooled together. Within markets, parameters are estimated at the level of the occupational skill type × job class $k = 1, ..., 5$, where job class is determined using the approach of Bonhomme et al. (2019). The sample is the subset of the BTS-Postes 2017-2023 (excluding 2020) that is matched to the Labor Force Survey (EEC) at the job level. Standard errors are clustered at the firm level.

5.5. How do supply and demand parameters determine wages and amenities?

To understand how different structural forces contribute to wage inequality, Table 8 decomposes wage variation into its constituent parts. The decomposition document how the estimated parameters of the model contribute to variation in observed wages. Panel A shows that three components exhibit comparable variance: residual wages (106% of total variance), amenity costs (101%), and productivity (85%). These shares exceed 100% because the components are negatively correlated with each other—when one is high, others tend to be low—which compresses the total wage variance relative to the sum of individual variances. If they were independent, these components would generate nearly three times the observed wage variance.

However, these variance shares alone do not reveal how the forces interact to shape wage inequality. Panel B presents correlations that illuminate three key economic mechanisms. First, productivity and amenity costs are strongly negatively correlated, attenuating wage variance by over 111 percent of the final magnitude. This correlation reflects that high-productivity firms face high costs to provide desirable flexibility. This finding is a combination of two features: the cost parameter is larger for high-paying jobs, and equilibrium amenities are higher because these firms tend to be in more competitive labor markets. Second, demanding jobs receive compensating wage differentials in *residual* wages. The negative correlation between amenity costs and residuals indicates that jobs

TABLE 8. Structural Sources of Wage Variation: Variance and Key Correlations

Panel A: Variance Contributions

Component	**Term**	**Variance**	**Share (%)**
Competition (markdown)	$\log \frac{\varepsilon^w_{oj}}{1+\varepsilon^w_{oj}}$	0.002	0.9
Firm Productivity	ψ_{jt}	0.209	85.1
Amenity Cost	$\rho_j \cdot \delta_{oj} \cdot a_{ojt}$	0.248	100.8
Skill Premium	$\log \theta_{ol}$	0.068	27.6
Within-Job Returns to Scale	$\rho_j \cdot \log L_{ojt}$	0.022	9.1
Residual	ϵ^w_{ojt}	0.262	106.5

Panel B: Key Covariance Contributions

Relationship	**Term**	**Covariance**	**Share (%)**
Productivity × Amenity Cost	$\mathrm{Cov}(\psi_{jt}, \rho_j \cdot \delta_{oj} \cdot a_{ojt})$	–0.274	–111.5
Amenity Cost × Residual	$\mathrm{Cov}(\epsilon^w_{ojt}, \rho_j \cdot \delta_{oj} \cdot a_{ojt})$	–0.259	–105.6
Productivity × Residual	$\mathrm{Cov}(\psi_{jt}, \epsilon^w_{ojt})$	–0.130	–53.1
Skill Premium × Residual	$\mathrm{Cov}(\epsilon^w_{ojt}, \log \theta_{ol})$	0.146	59.6
Other Covariances		-0.047	-19.3
Total variance		**0.245**	**100.0**

Notes: Panel A reports the variance of each component of the structural wage equation (14). "Share" column indicates the share of the wage variance explained by the component. Shares exceed 100% because components are correlated; total variance equals 0.246. Panel B reports covariances between key components.

where amenities are costly to provide command higher wages beyond what the model predicts. Similarly, high-productivity firms pay somewhat lower residual wages, though this effect is smaller in magnitude. These patterns suggest that workers extract compensating premia for demanding work conditions both through the direct technological channel ($\rho \cdot \delta \cdot a$) and as unobserved wage components. Third, skill premia amplify wage inequality through their interaction with unobserved factors. High-skill workers (high θ_{ol}) earn substantially higher residual wages, suggesting returns to unobserved ability or match quality beyond measured skill differences.

TABLE 9. Structural Sources of Amenity Variation: Variance and Covariance

Panel A: Variance Contributions			
Component	**Term**	**Variance**	**Share (%)**
Relative preference for amenities	$\log \frac{\varepsilon^a_{oj}}{1+\varepsilon^w_{oj}}$	0.002	0.2
Amenity Cost	$\log(-\delta_{oj})$	0.253	22.3
Residual	ϵ^a_{ojt}	0.882	78.3
Panel B: Covariance Contributions			
Relationship	**Term**	**Covariance**	**Share (%)**
Preference × Amenity Cost	$\mathrm{Cov}(\log \frac{\varepsilon^a_{oj}}{1+\varepsilon^w_{oj}}, -\log(-\delta_{oj})$	−0.008	−0.7
Amenity Cost × Residual	$\mathrm{Cov}(\epsilon^a_{ojt}, -\log(-\delta_{oj})$	0.007	0.6
Preference × Residual	$\mathrm{Cov}(\psi_{jt}, \epsilon^a_{ojt})$	−0.002	−0.2
Total variance		**1.132**	**100.0**

Notes: Panel A reports the variance of each component of the structural amenity equation (15). "Share" column indicates the share of the amenity variance explained by the component. Shares exceed 100% because components are correlated; total variance equals 1.132. Panel B reports covariances between key components.

Table 9 provides a similar decomposition of the components of the structural amenity equation (15). The decomposition reveals that technology is central for amenities: in Panel B, the technology term explains about 23 percent of the variance in the amenity, while the term capturing workers' preferences for amenities relative to amenities is close to zero. This suggests that firm technology is a first-order driver of amenity provision, and worker preferences play a limited role. Additionally, almost no variation in the amenity is driven by covariances across the different terms.

Together, these decompositions illuminate the mechanisms generating observed compensating differentials and their implications for inequality. The wage decomposition reveals that high-productivity firms face high amenity costs, and these demanding jobs pay compensating wage

premia. The amenity decomposition shows this pattern is driven by technological constraints rather than worker sorting: technology costs account for 22% of amenity variance while worker preferences contribute only 0.2%. High-productivity environments technologically require demanding conditions and workers receive wage compensation for these constraints.

6. Equilibrium counterfactuals

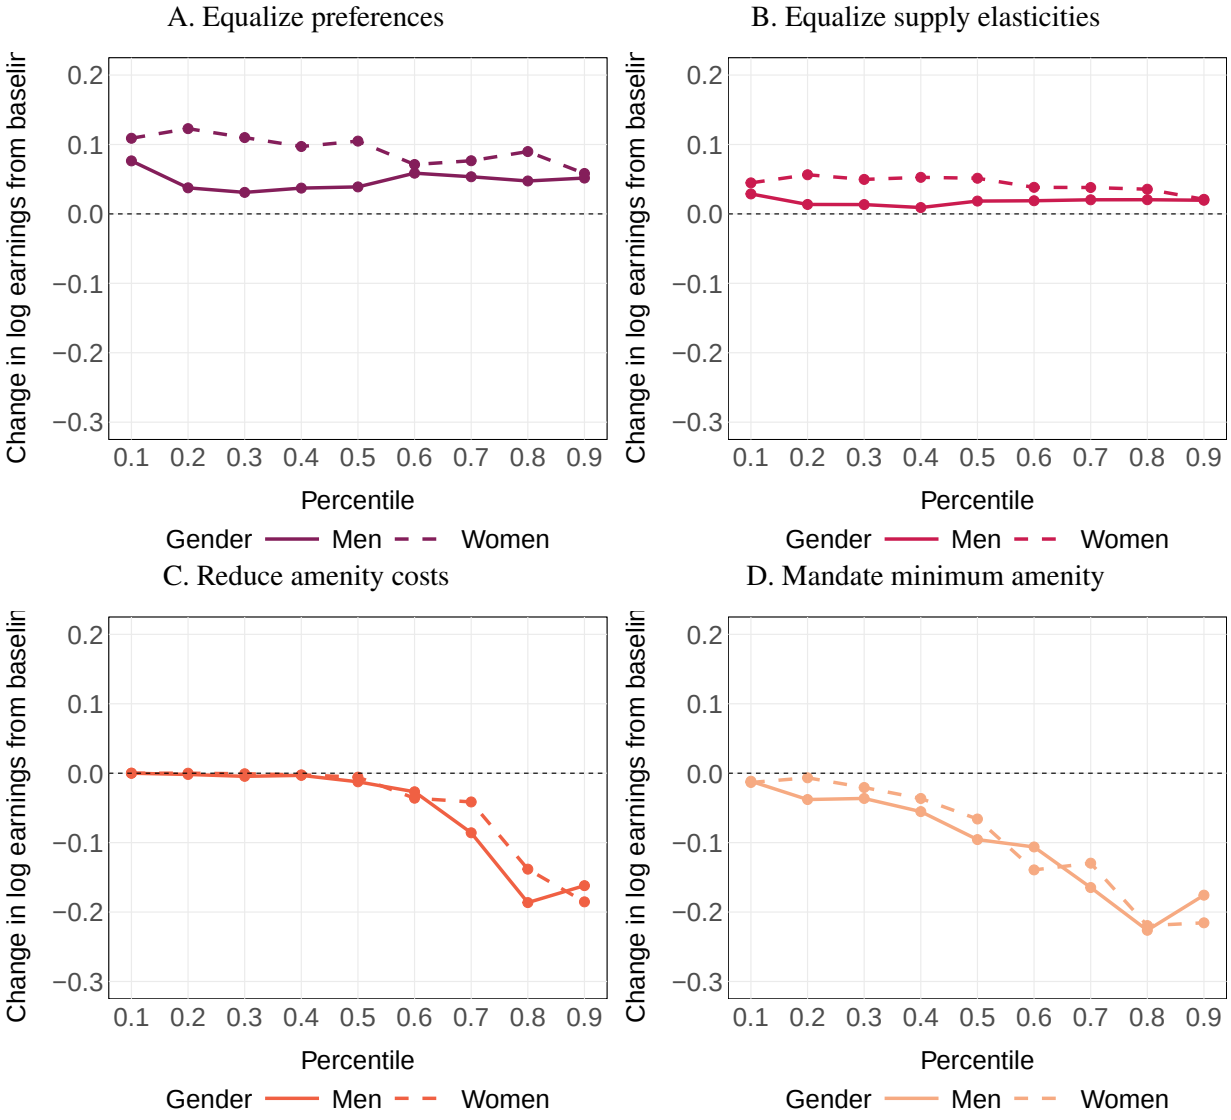
In this section, I simulate outcomes of counterfactual changes to model parameters. These counterfactuals allow me to quantify the contributions of firm costs and worker preferences to the determination of equilibrium amenities, wages, and gender inequality. In addition, I examine the impacts of mandating a minimum level of flexibility—analogous to the right-to-disconnect policies that are becoming increasingly common in Europe—for these same outcomes. In all counterfactuals, I abstract away from endogenous changes to labor supply elasticities that result from workers' resorting across jobs. In practice, because labor markets are large and workers have many alternatives, changes to elasticities due to marginal changes in labor market concentration are likely to be very small.

Equalizing preference parameters. I examine the impact of women's differences in preferences on the equilibrium by changing their utility coefficients α_f (on wages) and β_f (on flexibility) to be the same as men's. α_f increases from 0.55 to 0.80 and β_f increases from 0.71 to 0.76. The changes increase women's labor supply elasticities to all jobs, holding fixed the component of the elasticity that is determined by monopsony power within the choice set.

Panel A of Figure 12 shows that equalizing preferences raises women's equilibrium wages by 7.7 on average, with the largest pay gains for women in the bottom half of the wage distribution. Equilibrium wages for men also rise by 3-8 percent across the distribution. Intuitively, because women's preferences tilt away from amenities and towards wages, equilibrium amenity provision falls for both women and men. Firms offset this reduction in amenities by increasing men's wages. Panel A of Figure 13 shows that the reduction in amenity cost is more than offset by wage increases for both genders. This result highlights the role of monopsony power in the labor market in the determination of *both* wages and amenities in equilibrium.

Equalizing preferences also reduces the total gender wage gap by 4.6 percentage points, from 22.9 percent to 18.3 percent, attributing 20 percent of the baseline gap to these preference differences. The within-job gender gap declines by 1.2 percentage points, from 11 percent to 9.8 percent. Eequalizing preferences increases women's labor supply elasticities at each job, narrowing the job-specific markdown gap. After the preference adjustment, women sort towards better-paying firms, narrowing the gender gap attributed to sorting from 11.8 percent to 8.5 percent. Importantly,

FIGURE 12. Counterfactual changes to log wage distributions vs. baseline by gender



Notes: This figure plots differences in log earnings of men's and women's wage distributions under each of four counterfactuals relative to the observed baseline. I calculate deciles of the wage distribution under each counterfactual and compare to the same decile in the baseline wage distribution to obtain percentage changes in wages across the distribution. In each panel, differences in the wage distribution for men are plotted using solid lines and differences for women are plotted using dotted lines. Panel A shows results from the counterfactual that sets women's preference parameters in the utility equation (coefficients on wages and amenities) equal to men's. Panel B shows results from setting women's labor supply elasticities (with respect to both wages and amenities) in each job equal to men's in the same job. Panel C shows results from reducing the cost δ_{ok} to provide flexibility to workers in high-skill occupations. Panel D shows results from setting a minimum amenity (flexibility) level at the 20th percentile of the ex ante amenity distribution. The sample is the subset of the 2018 file of the BTS-Postes that is matched to the Labor Force Survey (EEC) at the job level.

equalizing preferences does not eliminate the gender wage gap entirely: men and women still earn different wages because they face different choice sets, markdowns, and skill premia.

Equalizing labor supply elasticities. I next set women's preference parameters and labor supply elasticities equal to men's at each job. This adds to the first counterfactual by additionally equalizing the competitive labor market conditions faced by men and women under monopsonistic competition.

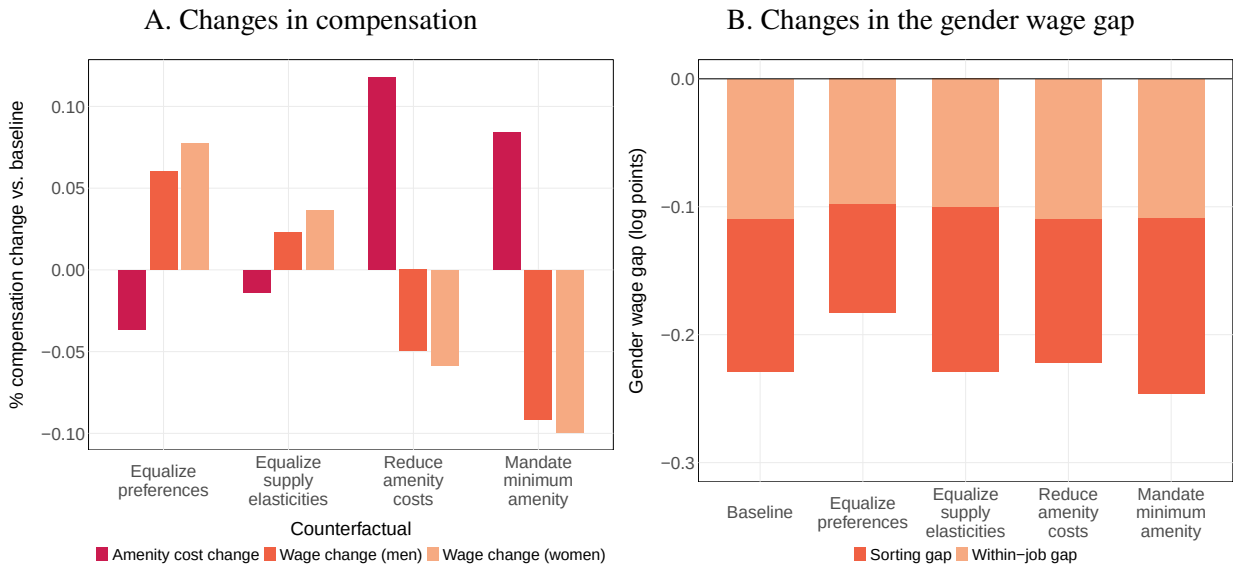
Wage gains are more modest under equalized elasticities than under only equalized preferences: across the wage distribution, women experience gains of 2-6 percent and men experience gains of 1-2 percent. Intuitively, this difference results from the slightly higher equilibrium amenities that are provided under this counterfactual relative to equalized preferences only. The comparison shows that equilibrium amenity provision is somewhat higher when labor markets are more competitive. Additionally, the equalizing elasticities has virtually no impact on the gender gap, suggesting that equalizing differences in imperfect competition offsets the reduction in sorting caused by equalizing preferences. Comparing these two counterfactuals sheds light on an important, but subtle, feature of the labor market: amenity provision is partly driven by imperfect competition.

Reducing the cost of flexibility for high-skill workers. What would happen to wages, amenities, gender gaps, and worker welfare if firms faced lower costs to provide flexibility to high-skill workers? To answer this question, I reduce the cost parameter δ_{ok} for the highest tier of occupational skill (1-digit occupation code 3) to equalize it to the cost parameter for the second skill tier (1-digit occupation code 4) in the same job type. For example, a job of type $k = 1$ will set $\delta_{3,1} = \delta_{4,1}$. The parameters δ_{ok} for all other occupational skill types k remain unchanged.

Panel C of Figure 12 shows that wages in the top half of the wage distribution decline for both men and women. As amenities become relatively cheaper to provide, firms shift compensation towards amenities and away from wages. Panel A of Figure 13 shows that firms spend more in total on compensation: the average decrease in wages for both men and women is more than offset by an increase in amenity spending. This happens because the decrease in input costs increases optimal output levels, thereby raising marginal product. This outweighs the direct effect of a larger negative $\rho_j \cdot \delta_{oj} \cdot a_{ojt}$ pushing equilibrium wages downward. Panel A of Figure 14 illustrates that this increase in compensation translates into welfare gains of 1-3 percent for workers across the distribution. Additionally, the gender wage gap narrows modestly—by 0.7 percentage points—driven by a decline in the sorting gap. Intuitively, the cost reduction also narrows differences in amenities across firms, which dampens sorting effects.

Mandating a minimum amenity level. How would obligating firms to provide a minimum level of flexibility affect workers in equilibrium? In this counterfactual, I require firms to provide the flexibility amenity at the 20th percentile of the baseline distribution. Because no model parameters change, jobs whose equilibrium amenities exceeded this minimum before the policy will not

FIGURE 13. Counterfactual changes to compensation and gender gaps



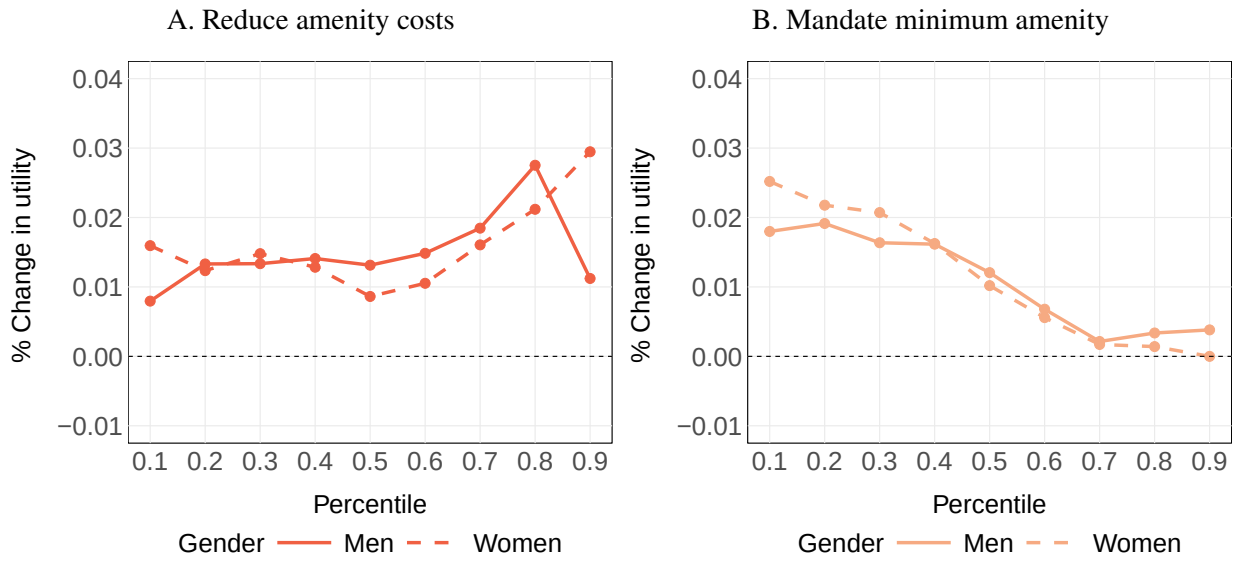
Notes: This figure plots changes to the means of estimated determinants of compensation (Panel A) and the gender wage gap (Panel B) under each counterfactual. In Panel A, the "amenity cost change" bar reflects the change in mean total spending on amenities per worker. In Panel B, the total gap is the difference in log wage, and the within-job gap is the difference in log wage after controlling for job fixed effects. The sorting gap is the share of the total gap that is not explained by the within-job gap. The sample is the subset of the 2018 file of the BTS-Postes that is matched to the Labor Force Survey (EEC) at the job level.

change their amenities.[13] This counterfactual has close analogues in contemporary policy debates, especially in Europe. Many countries have considered or established policies providing for a "right to disconnect," which protects workers from punishment for disconnecting from their work outside of regular business hours. For example, under such a policy, workers cannot be punished for not responding to e-mails after the end of the business day. France was one of the first countries to recognize such a right, beginning with case law in the early 2000s (Cour de cassation, chambre sociale 2004) and introduced into the French Labor Code in 2016 (République Française 2021). The current version of the law obligates firms to establish policies, but leaves them to set the terms by negotiations between firms and unions.

Panel D of Figure 12 shows that the policy induces firms to reduce wages to substitute towards amenities, following a similar pattern to the cost reduction counterfactual. However, in this scenario, the wage reductions are larger across most of the wage distribution and extend to jobs that had lower pay before the policy. Jobs that offered low levels of flexibility—mostly those in the upper third of

[13]This result relies on the simplifying assumption that labor supply elasticities to individual firms do not change. Because labor markets are large and dense, reallocation across jobs under the counterfactual results in minimal changes to the distribution of market shares across jobs within a market. Allowing for changes in elasticities would result in small endogenous changes to amenities for jobs for which the mandate does not bind, but because reallocation is minimal these changes would be small.

FIGURE 14. Counterfactual changes to worker welfare vs. baseline by gender



Notes: This figure plots changes in deciles of the estimated distribution of worker utility for men and women relative to the baseline under two counterfactuals: reducing amenity costs (Panel A) and mandating a minimum amenity at the 20th percentile of the ex ante distribution (Panel B). Changes in utility are calculated using baseline preferences under equation 7. The sample is the subset of the 2018 file of the BTS-Postes that is matched to the Labor Force Survey (EEC) at the job level.

the wage distribution—cut wages to compensate for the cost of complying with the mandate. Panel A of Figure 13 shows that, in contrast to the cost reduction counterfactual, the increase in spending on amenities is *more* than offset by a reduction in wages on average, reducing total compensation. Gender wage gaps also increase by 1.8 percentage points, driven by an increase in the sorting gap. Because firms with high amenity costs cut wages more, they become disproportionately more attractive to women, increasing gendered sorting in equilibrium.

The welfare incidence of the mandate reflects heterogeneity in labor market competition throughout the economy. High-skill, high-pay jobs with high costs to provide flexibility are concentrated in the dense, highly competitive Paris-area labor market. Requiring firms in this market to comply with restrictive amenity mandates leads them to trade off amenity costs with wage costs, so that workers' welfare remains nearly unchanged. In contrast, in less competitive markets, requiring minimum amenities has less of an impact on wages and induces firms to provide better total compensation. As a result, lower-quality jobs deliver higher utility under the mandate; worker welfare in these jobs increases by 1-2 percent. This finding is analogous to results from the minimum wage literature that regulating pay or benefits can be welfare-improving when firms have monopsony power. In contrast, in high-pay, high-skill markets, the amenity increase is entirely offset by wage decreases.

7. Conclusion

This paper identifies that firms' production technologies play a crucial role in determining workplace amenities, particularly scheduling flexibility and demands on workers' time. Descriptive evidence reveals patterns that compensating differentials alone cannot explain: firms set amenities systematically across occupations, and high-paying firms offer both better and worse amenities rather than uniformly poor non-wage benefits. By developing and estimating an equilibrium model that incorporates both worker preferences and technological constraints, I uncover a sharp skill gradient in the productivity costs of providing flexibility: high-skill occupations face productivity penalties roughly twice as large as mid-skill occupations when offering favorable scheduling amenities.

Counterfactual simulations reveal that preferences and technology jointly determine labor market outcomes in ways that create difficult policy trade-offs. Equalizing women's amenity preferences to men's would reduce the gender wage gap by approximately 20 percent. However, policies mandating minimum flexibility levels—similar to Europe's expanding "right to disconnect" laws—would improve welfare for workers in less competitive markets but widen the gender wage gap by increasing sorting. This unintended consequence highlights how well-intentioned flexibility mandates may exacerbate the very inequalities they aim to address.

These findings reshape our understanding of workplace inequality and the scope for policy intervention. The concentration of inflexible jobs in high-productivity, high-skill occupations reflects fundamental technological constraints rather than market failures. Workers facing demanding schedules typically have outside options with similar limitations, creating market-level rigidities that individual firm interventions cannot easily overcome. The evidence suggests that achieving meaningful improvements in job quality and reducing gender inequality will require not just changing preferences or mandating standards, but changing production technologies.

Several directions for future research emerge from this framework. First, understanding how firms adapt their production technologies over time—distinguishing between what is inherently difficult to change versus what is malleable—remains an open question. The introduction of remote work provides a particularly compelling natural experiment. The COVID-19 pandemic simultaneously introduced a technological shock (organizational learning through mass adoption of remote work technology) and a preference shock (shifting social norms around workplace flexibility). The framework developed in this paper provides a starting point for disentangling supply and demand forces in the new equilibrium. Future work can teach us about whether technological constraints can evolve when firms are forced to experiment with new production methods.

Second, future work could design targeted interventions to understand the impact of reducing technological barriers to flexibility. Collaborating with large firms to experimentally vary production

processes or scheduling systems could provide causal evidence on which technological constraints are binding and which can be relaxed through organizational innovation. Such field experiments could also reveal whether firms underinvest in flexibility-enabling technologies due to coordination failures, information frictions, or organizational inertia. As labor markets continue evolving and workers increasingly value non-wage characteristics, understanding the role of technological change in amenity provision will be essential for designing effective policies that improve worker welfare.

References

Abowd, John M., Francis Kramarz, and David N. Margolis, "High Wage Workers and High Wage Firms," *Econometrica*, 1999, *67* (2), 251–333. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/1468-0262.00020.

Azar, José and Ioana Marinescu, "Monopsony Power in the Labor Market," in "Handbook of Labor Economics" 2024.

___ , Steven Berry, and Ioana Marinescu, "Estimating Labor Market Power," 2022.

Bai, Jushan and Serena Ng, "Determining the Number of Factors in Approximate Factor Models," *Econometrica*, 2002, *70* (1), 191–221.

Barbanchon, Thomas Le, Roland Rathelot, and Alexandra Roulet, "Gender Differences in Job Search: Trading off Commute against Wage," *The Quarterly Journal of Economics*, 2021, *136* (1), 381–426.

Bell, Alex, "Job Amenities and Earnings Inequality," June 2025.

Berger, David, Kyle Herkenhoff, and Simon Mongey, "Labor market power," *The American Economic Review*, 2022, *112* (4), 1147–1193.

Bloom, Nicholas, James Liang, John Roberts, and Zhichun Jenny Ying, "Does Working from Home Work? Evidence from a Chinese Experiment," *The Quarterly Journal of Economics*, February 2015, *130* (1), 165–218.

___ , Ruobing Han, and James Liang, "How Hybrid Working From Home Works Out," July 2022.

Bonhomme, Stéphane and Grégory Jolivet, "The Pervasive Absence of Compensating Differentials," *Journal of Applied Econometrics*, 2009, *24*, 763–795.

___ , Thibaut Lamadon, and Elena Manresa, "A Distributional Framework for Matched Employer-Employee Data," *Econometrica*, 2019, *87* (3), 699–739.

___ , ___ , and ___ , "Discretizing Unobserved Heterogeneity," *Econometrica*, 2022, *90* (2), 625–643.

Brown, Charles, "Equalizing Differences in the Labor Market," *The Quarterly Journal of Economics*, 1980, *94* (1), 113–134. Publisher: Oxford University Press.

Burdett, Kenneth and Dale T. Mortensen, "Wage Differentials, Employer Size, and Unemployment," *International Economic Review*, 1998, *39* (2), 257–273. Publisher: [Economics Department of the University of Pennsylvania, Wiley, Institute of Social and Economic Research, Osaka University].

Caldwell, Sydnee, Ingrid Haegele, and Jörg Heining, "Firm Pay and Worker Search," 2025.

Card, David, "Who set your wage?," *The American Economic Review*, 2022, *112* (4), 1075–1090.

___ , Ana Rute Cardoso, Jörg Heining, and Patrick Kline, "Firms and Labor Market Inequality: Evidence and Some Theory," *Journal of Labor Economics*, 2018, *36* (51).

Chen, M. Keith, Peter E. Rossi, Judith Chevalier, and Emily Oehlsen, "The Value of Flexible Work: Evidence from Uber Drivers," *Journal of Political Economy*, 2019, *127* (6), 2735–2794.

Collis, Manuela and Clémentine Van Effenterre, "Workplace Hostility."

Corradini, Viola, Lorenzo Lagos, and Garima Sharma, "Collective Bargaining for Women: How Unions Can Create Female-Friendly Jobs," *The Quarterly Journal of Economics*, May 2025, p. qjaf024.

Cortes, Guido Matias and Giovanni Gallipoli, "The costs of occupational mobility: An aggregate analysis," *Journal of the European Economic Association*, 2018, *16* (2), 275–315.

Cour de cassation, chambre sociale, "Arrêt n°01-45.889," February 2004.

Cubas, German, Chinhui Juhn, and Pedro Silos, "Work-Care Balance over the Day and the Gender Wage Gap," *AEA Papers and Proceedings*, 2021, *111*, 149–153.

___ , ___ , and ___ , "Coordinated Work Schedules and the Gender Wage Gap," *The Economic Journal*, 2023, *133* (651), 1036–1066.

Danon, Leon, Albert Díaz-Guilera, Jordi Duch, and Alex Arenas, "Comparing community structure identification," *Journal of Statistical Mechanics*, 2005, *9*.

Dube, Arindrajit, Suresh Naidu, and Adam D Reich, "Power and Dignity in the Low-Wage Labor Market: Theory and Evidence from Wal-Mart Workers."

Eaton, Jonathan and Samuel Kortum, "Technology, Geography, and Trade," *Econometrica*, 2002, *70* (5), 1741–1779.

Emanuel, Natalia and Emma Harrington, "Working Remotely? Selection, Treatment, and the Market for Remote Work," *American Economic Journal: Applied Economics*, October 2024, *16* (4), 528–559.

___ , ___ , and Amanda Pallais, "The Power of Proximity to Coworkers: Training for Tomorrow or Productivity Today?"

Eriksson, Tor and Nicolai Kristensen, "Wages or Fringes? Some Evidence on Trade-Offs and Sorting," *Journal of Labor Economics*, 2014, *32* (4), 899–928. Publisher: [The University of Chicago Press, Society of Labor Economists, NORC at the University of Chicago].

Fan, Ying, "Ownership Consolidation and Product Characteristics: A Study of the US Daily Newspaper Market," *American Economic Review*, 2013, *103* (5), 1598–1628.

Felix, Mayara, "Trade, Labor Market Concentration, and Wages," 2022.

Fogel, Jamie and Bernardo Modenesi, "What is a Labor Market? Classifying Workers and Jobs Using Network Theory," 2022.

Garro-Marín, César, Neil Thakral, and Linh T. Tô, "Work Hours and Amenity Trade-Offs," *AEA Papers and Proceedings*, May 2025, *115*, 282–286.

Goldin, Claudia, "A Grand Gender Convergence: Its Last Chapter," *The American Economic Review*, 2014, *104* (4), 1091–1119.

___ and Lawrence F. Katz, "The Cost of Workplace Flexibility for High-Powered Professionals," *The Annals of the American Academy of Political and Social Science*, 2011, *638*, 45–67. Publisher: [Sage Publications, Inc., American Academy of Political and Social Science].

Gruber, Jonathan and Alan Krueger, "he Incidence of Mandated Employer- Provided Insurance: Lessons from Workers' Compensation Insurance," *Tax Policy and the Economy*, 1991, *5*, 111–143.

Hall, Robert E. and Andreas I. Mueller, "Wage Dispersion and Search Behavior: The Importance of Nonwage Job Values," *Journal of Political Economy*, August 2018, *126* (4), 1594–1637. Publisher: The University of Chicago Press.

Hamermesh, Daniel S., "The Timing of Work over Time," *The Economic Journal*, 1999, *109* (452), 37–66. Publisher: [Royal Economic Society, Wiley].

Hazell, Jonathon, Christina Patterson, Heather Sarsons, and Bledi Taska, "National Wage Setting," 2022.

Holland, Paul W., Kathryn Blackmond Laskey, and Samuel Leinhardt, "Stochastic Block Models: First Steps," *Social Networks*, 1983, *5* (2), 109–137.

Humlum, Anders, Mette Rasmussen, and Evan Rose, "Firm Premia and Match Effects in Pay vs. Amenities," May 2025.

Hwang, Hae-shin, Dale T. Mortensen, and W. Robert Reed, "Hedonic Wages and Labor Market Search," *Journal of Labor Economics*, October 1998, *16* (4), 815–847. Publisher: The University of Chicago Press.

Jarosch, Gregor, Jan Nimczik, and Isaac Sorkin, "Granular Search, Market Structure, and Wages," *Review of Economic Studies*, 2024, *91* (6), 3569–3607.

Juhn, Chinhui and Kristin McCue, "Specialization Then and Now: Marriage, Children, and the Gender Earnings Gap across Cohorts," *Journal of Economic Perspectives*, 2017, *31* (1), 183–204.

Karrer, Brian and M.E.J. Newman, "Stochastic blockmodels and community structure in networks," *Physical Review*, 2011, *83* (1).

Kim, Hyuncheol Bryant, Hyunseob Kim, and John Y. Zhu, "The Selection Effects of Part-Time Work: Experimental Evidence from a Large-Scale Recruitment Drive," 2024.

Kline, Patrick, "Labor Market Monopsony: Fundamentals and Frontiers," 2025.

Kroft, Kory, Yao Luo, Magne Mogstad, and Bradley Setzler, "Imperfect Competition and Rents in Labor and Product Markets: The Case of the Construction Industry," *American Economic Review*, 2025, *115* (9).

Lachowska, Marta, Alexandre Mas, Raffaele Saggio, and Stephen Woodbury, "Work Hours Mismatch," 2023.

Lagos, Lorenzo, "Union Bargaining Power and the Amenity-Wage Tradeoff," 2025.

Lamadon, Thibaut, Jeremy Lise, Costas Meghir, and Jean-Marc Robin, "Labor Market Matching, Wages, and Amenities," 2024.

___ **, Magne Mogstad, and Bradley Setzler**, "Imperfect Competition, Compensating Differentials, and Rent Sharing in the US Labor Market," *American Economic Review*, January 2022, *112* (1), 169–212.

Lang, Kevin and Sumon Majumdar, "The Pricing of Job Characteristics When Markets Do Not Clear: Theory and Policy Implications," *International Economic Review*, 2004, *45* (4), 1111–1128. Publisher: [Economics Department of the University of Pennsylvania, Wiley, Institute of Social and Economic

Research, Osaka University].

Lavetti, Kurt, "Compensating Wage Differentials in Labor Markets: Empirical Challenges and Applications," *Journal of Economic Perspectives*, August 2023, *37* (3), 189–212.

___ **and Ian M Schmutte**, "Estimating Compensating Wage Differentials with Endogenous Job Mobility," 2025.

Lehmann, Tobias, "Non-Wage Job Values and Implications for Inequality," 2025.

Maestas, Nicole, Kathleen J. Mullen, David Powell, Till Von Wachter, and Jeffrey B. Wenger, "The Value of Working Conditions in the United States and Implications for the Structure of Wages," *American Economic Review*, July 2023, *113* (7), 2007–2047.

Manning, Alan, "Monopsony in labor markets: a review," *ILR Review*, 2021, *74* (1), 3–26.

Mas, Alexandre, "Non-Wage Amenities," in "Handbook of Labor Economics," 1 ed., Vol. 6 2025.

___ **and Amanda Pallais**, "Valuing Alternative Work Arrangements," *American Economic Review*, December 2017, *107* (12), 3722–3759.

___ **and** ___ , "Labor Supply and the Value of Non-Work Time: Experimental Estimates from the Field," *American Economic Review: Insights*, June 2019, *1* (1), 111–126.

___ **and** ___ , "Alternative Work Arrangements," *Annual Review of Economics*, August 2020, *12*, 631–658.

Morchio, Iacopo and Christian Moser, "The Gender Pay Gap: Micro Sources and Macro Consequences," 2025.

Nagler, Markus, Johannes Rincke, and Erwin Winkler, "Working from home, commuting, and gender," *Journal of Population Economics*, *37* (58).

Nimczik, Jan, "Job Mobility Networks and Data-driven Labor Markets," 2023.

Ommeren, Jos Van, Gerard J. Van Den Berg, and Cees Gorter, "Estimating the Marginal Willingness to Pay for Commuting," *Journal of Regional Science*, 2000, *40* (3), 541–563. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/0022-4146.00187.

Ouimet, Paige and Geoffrey Tate, "Firms with Benefits? Nonwage Compensation and Implications for Firms and Labor Markets."

Oyer, Paul, "Salary or Benefits?," in "Work, Earnings and Other Aspects of the Employment Relation," Vol. 28, Emerald Group Publishing Limited, 2008, pp. 429–467.

Pierce, Brooks, "Compensation Inequality," *The Quarterly Journal of Economics*, 2001, *116* (4), 1493–1525.

Rosen, Sherwin, "Hedonic Prices and Implicit Markets: Product Differentiation in Pure Competition," *Journal of Political Economy*, January 1974, *82* (1), 34–55.

___ , "The theory of equalizing differences," in "Handbook of Labor Economics," Vol. 1, Elsevier, 1986, pp. 641–692.

Roussille, Nina and Benjamin Scuderi, "Bidding for Talent: A Test of Conduct in a High-Wage Labor

Market."

République Française, "Article L2242-17," 2021. Accessed: October 24, 2025.

Sockin, Jason, "Show Me the Amenity: Are Higher-Paying Firms Better All Around?," 2022.

Sokolova, Anna and Todd Sorensen, "Monopsony in labor markets: a meta-analysis," *ILR Review*, 2021, *74* (1), 27–55.

Sorkin, Isaac, "Ranking Firms Using Revealed Preference," *The Quarterly Journal of Economics*, August 2018, *133* (3), 1331–1393.

Sullivan, Paul and Ted To, "Search and Nonwage Job Characteristics," *The Journal of Human Resources*, 2014, *49* (2), 472–507. Publisher: [University of Wisconsin Press, Board of Regents of the University of Wisconsin System].

Summers, Lawrence H., "Some Simple Economics of Mandated Benefits," *The American Economic Review*, 1989, *79* (2,), 177–183.

Taber, Christopher and Rune Vejlin, "Estimation of a roy/search/compensating differential model of the labor market," *Econometrica*, 2020, *88* (3), 1031–1069.

Villanueva, Ernesto, "Estimating Compensating Wage Differentials Using Voluntary Job Changes: Evidence from Germany," *ILR Review*, July 2007, *60* (4), 544–561. Publisher: SAGE Publications Inc.

Volpe, Oscar, "Job Preferences, Labor Market Power, and Inequality," 2025.

Wiswall, Matthew and Basit Zafar, "Preference for the Workplace, Investment in Human Capital, and Gender," *The Quarterly Journal of Economics*, February 2018, *133* (1), 457–507.

Appendix A. Estimation Appendix

I cannot directly estimate a type for each worker given the limitations of the data.[14] However, I can use the approach of Bonhomme et al. (2019) (henceforth BLM) to discretize job and worker heterogeneity. The BLM algorithm has two steps. First, I classify jobs into latent types using k-means clustering on 20 quadrature points of the job's earnings distribution. Second, I estimate (i) the composition of worker types and (ii) the wage distribution for each worker type within each firm class. I refer the interested reader to the original paper for greater detail on the classification procedure. In this Appendix, I show the results from the baseline BLM classification and present an extension that allows me to identify job-specific type distributions and job type × worker type wages. I then use the extension to document gender heterogeneity in skill premia.

A.1. Baseline distributions of unobserved heterogeneity

Figure A.1 shows the baseline results characterizing the distributions of worker types across job types. Types are computed for two groups of markets: the Paris high-skill market and all other markets pooled together. Job types $k = 1, ..., 5$ and worker types $l = 1, ..., 5$ are ordered so that type 5 corresponds to the highest-paid type of each. In the figures, darker colors correspond to lower-paid worker types.

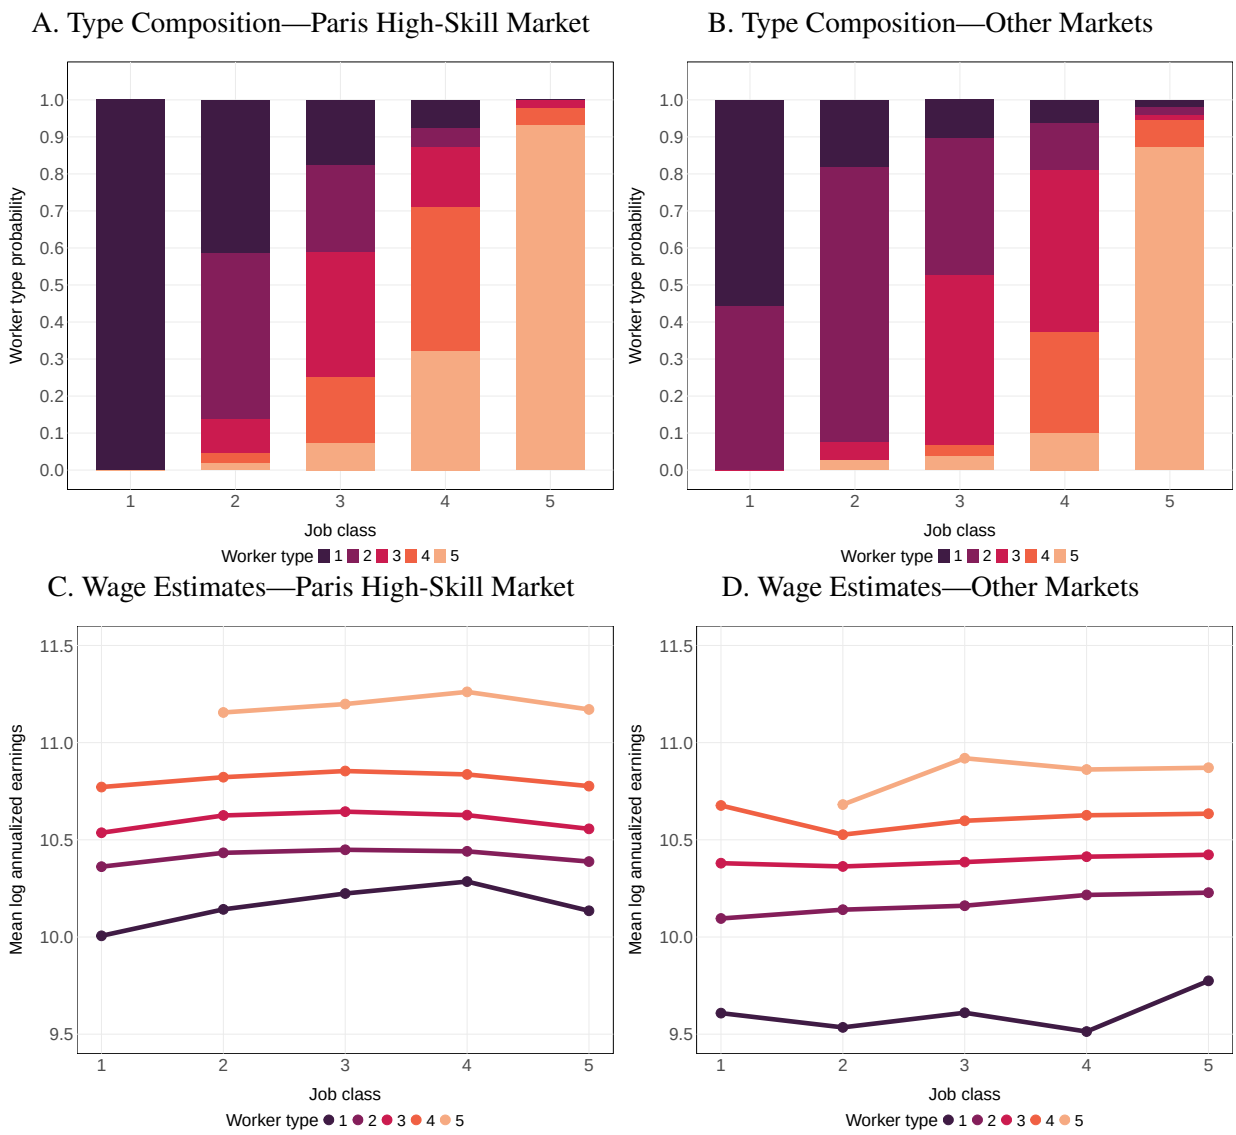
Panels A and B show that workers exhibit strong sorting across firms in both groups: the lowest-paying firm type consists almost entirely of the lowest-paid two worker types; the vast majority of workers in the highest-paying firm type belong to the highest worker type. While the middle firm types exhibit more type diversity, clear patterns of assortativity persist across the job type distribution. Panels C and D show minimal evidence of complementarities in worker pay across job types: workers of each type are paid similar wages regardless of the job type. Both of these sets of results are similar to the main findings of Bonhomme et al. (2019).

A.2. Estimation of job-specific type distributions with Bayesian hierarchical shrinkage

Within each job type k, the BLM approach identifies (i) the composition of worker types $p_k(l)$ and (ii) the wage distribution of each worker type $f_{kl}(w)$, assuming a Normal distribution. It does not identify the type distribution or wage distribution of types within specific *jobs* $j \in \mathcal{J}_k$. I estimate job-specific distributions of worker types and wages conditional on type using the following Bayesian hierarchical shrinkage algorithm.

[14] Specifically, the panel subset is an as-good-as-random 1/12 subset of the population, and I do not observe education level for the full population. Thus in order to obtain conditional wage distributions of worker and job types, I cannot use AKM or leverage education as an observed skill proxy.

FIGURE A.1. Histograms of Estimated Production Parameters



Notes: This figure plots the estimated distributions of job and worker types (Panels A and B) and the estimated mean wage for each worker × job type combination (Panels C and D) using the approach of Bonhomme et al. (2019). There are $K = 5$ job types and $L = 5$ worker types, where type 1 is the lowest-pay type and type 5 is the highest-pay type for each. Jobs are defined at the level of the enterprise × 2-digit occupation × département and are clustered into types using 20 quadrature points of the wage distribution. Job types are estimated on the 2018 file of the BTS-Postes and extrapolated to other years using level shifts.

a. Define the **class-level priors**: $p_k(l)$ is the distribution of worker types in firm class k and $f_{kl}(w)$ is the distribution of wages for type l in class k. We know both from the initial BLM estimation. Use $p_k(l)$ as the initial guess of $q_{oj}(l)$.
b. **E-step**: conditional on a guess q_{oj}, define τ_{il}, the probability that i is of type l given the observed wage y_i and type mix q_{oj}:

$$\tau_{il} = \frac{q_j(l) \cdot f_{kl}(w_i)}{\sum_{l'=1}^{L} q_j(l') \cdot f_{kl'}(w_i)}. \tag{A1}$$

Then define **soft counts** $\hat{n}_l = \sum_i \tau_{il}$ and $n_j = \sum_l \hat{n}_l$ (job size).
c. **M-step**: update the guess $q_{oj}(l)$ using the MAP update:

$$q_l^{new} = \frac{\lambda \cdot p_k(l) - 1 + \hat{n}_l}{\lambda - L + n_j}. \tag{A2}$$

d. Iterate through the E-step and M-step until convergence.

The type composition algorithm yields type probabilities q_l^* and an associated τ_{il}^* for each worker. I next apply an Empirical Bayes shrinkage estimator to estimate the mean wage for the job × worker type, $\mu_{l|oj}$. Define a Normal prior for $\mu_{l|oj}$ using the class mean μ_{kl} and standard error σ_{kl} obtained from the BLM estimation step. The posterior mean is the shrinkage estimator:

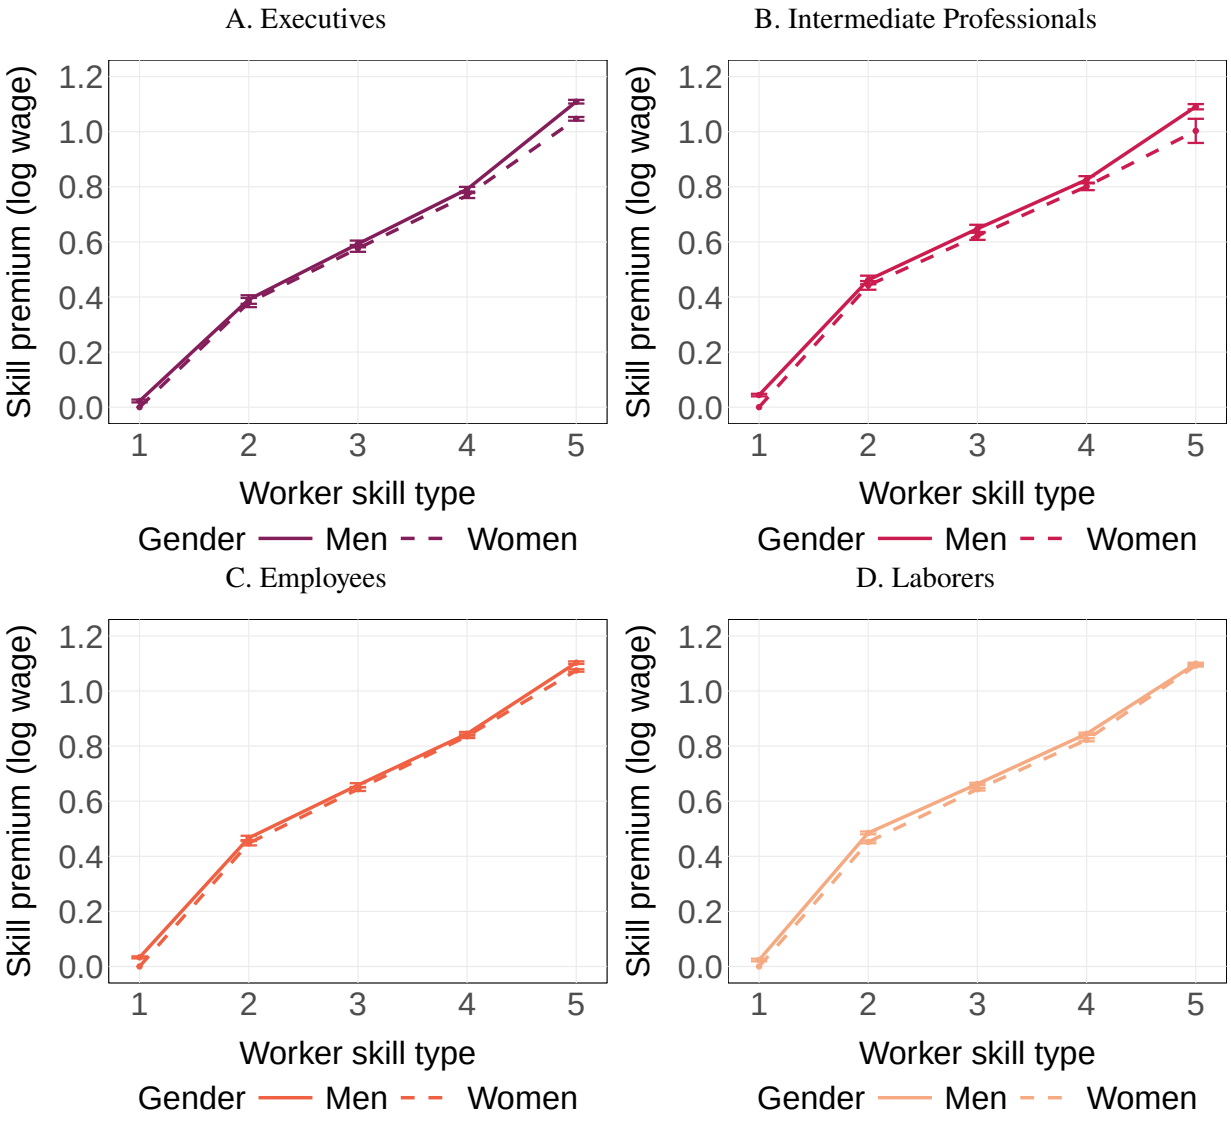
$$\hat{\mu}_{l|oj}^{EB} = \frac{\tau_{0,kl} \cdot \mu_{kl} + \sum_i \tau_{il}^* \cdot y_i / \sigma_{kl}^2}{\tau_{0,kl} + \sum_i \tau_{il}^* / \sigma_{kl}^2}, \tag{A3}$$

where $\tau_{0,kl}$ is the prior precision hyperparameter, defined as $\tau_{0,kl} \equiv n_0/\sigma_{kl}^2$. n_0 is a pseudo-sample size that governs how strongly I shrink the job-specific mean toward the class mean. Intuitively, if the effective number of workers in the job is much smaller than n_0, I shrink more towards the prior; if the effective number of workers is much larger than n_0, I rely on the data to a greater extent. In the empirical implementation I choose $n_0 = 10$.

A.3. Gender heterogeneity in skill premia

Figure A.2 shows estimates of skill premia $\log \theta_{olg}$ for men and women in each of the four major occupational categories. All skill premia are defined relative to women in the lowest skill group. Across the skill distribution, there are small bt precisely estimated differences in skill premia for men and women. The gaps are largest in the higher-skill occupations gap is 0.063 log points for executives and 0.087 log points for intermediate professionals. Gaps are smaller for low and intermediate skill types and for lower-skill occupations.

FIGURE A.2. Gender Differences in Skill Premia by Occupation



Notes: This figure plots estimated average skill premia $\log \theta_{olg}$ for occupational skill category o, worker skill type $l = 1, ..., 5$, and gender g = (men, women). Skill premia are estimated at the 2-digit occupation × gender × skill type level by integrating over the job-specific type distribution, identified using the algorithm in section A.2. The figures plot the means and standard errors of the means of these more detailed estimates. The sample is the subset of the BTS-Postes 2017-2023 (excluding 2020) that is matched to the Labor Force Survey (EEC) at the job level. Standard errors are clustered at the firm level.

Appendix B. Supplemental Exercises

B.1. Alternative sources of variation in job amenities

B.1.1. Do amenities vary within or between jobs?

To link job-level amenities to administrative data, I assume that the amenities I observe for individual workers in the large-scale surveys are representative of the amenities offered in their *job*, where the job is defined at the level of the firm × occupation × location. However, linking at this level of aggregation would sacrifice information and misspecify the model of amenity-setting if amenities vary substantially *within* jobs. To test this assumption and validate the linking approach, I decompose the variance in each measured amenity a into between- and within-firm variance:

$$Var(a_{ij}) = \underbrace{Var(\mathbb{E}[a_i|J = j])}_{\text{between firms}} + \underbrace{\mathbb{E}[Var(a_i|J = j)]}_{\text{within firms}},$$

where i represents an individual worker and j represents a job. I estimate the decomposition within-year to abstract from changes over time in amenity provision.

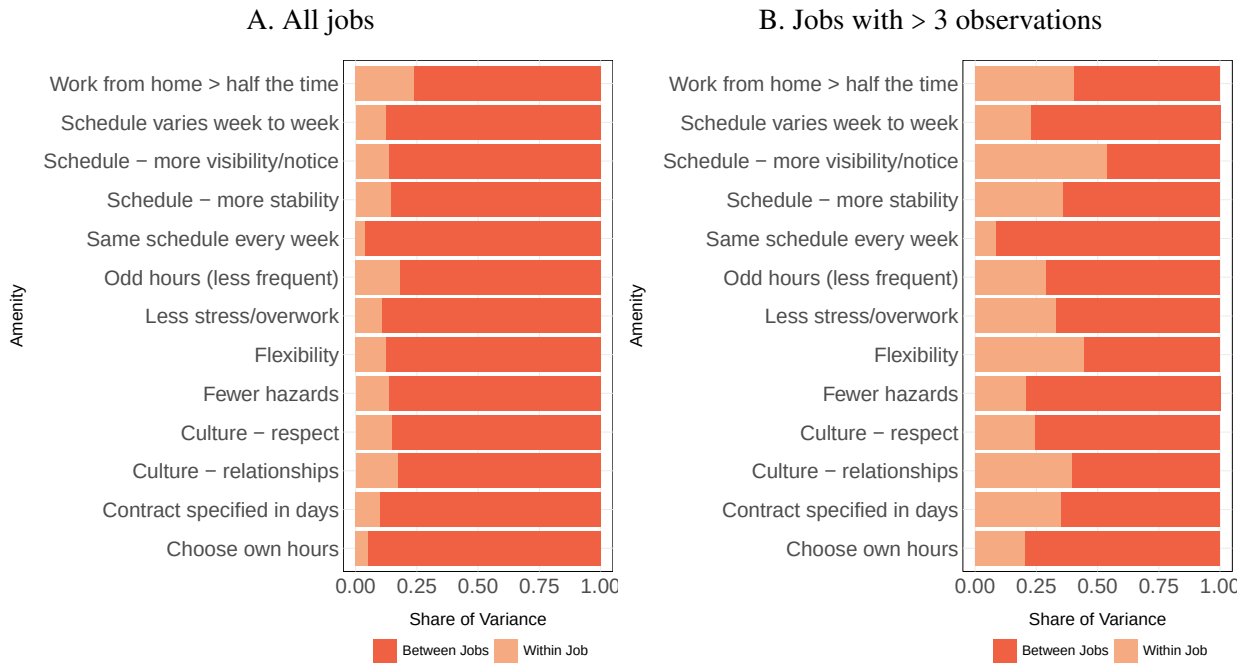
Figure B.1 shows the share of total variance in each amenity that is within-job (light orange) versus between jobs (dark orange). In panel (a), I include all job observations from the EEC and CT-RPS samples. This panel shows that for each measured amenity, the vast majority of the variance comes from differences between jobs rather than differences in reported amenities among workers in the same job.

One might be concerned that the high share of between-job variation could be driven by survey sampling, which leads us to observe only one worker in many jobs. To address this concern, in panel (b), I restrict the sample to jobs where I observe at least three workers in the survey sample. Even with this sample restriction, much more than half of the variation in most amenities occurs between jobs rather than within jobs. This suggests that sample properties are not driving the result and that most of the variation in amenities is between jobs.

B.1.2. Are amenities set by collective bargaining groups or by individual firms?

In France, sectoral collective bargaining agreements regulate many aspects of employment arrangements, including both wages and amenities. Collective bargaining has become increasingly decentralized over time; much of the negotiation occurs directly between unions and individual firms. However, one might be concerned that collectively bargained provisions restrict the amenities that firms can provide, leaving little room for firm-specific variation. If this were true, it would threaten the findings of Fact 1 and undermine the job- and firm-level variation in amenities that

FIGURE B.1. Within- and Between-Job Variation in Amenities



Notes: This figure shows the share of variance of each amenity explained by within-job versus between-job variation. The following amenities are derived from the Continuous Labor Force Survey (EEC) from 2016 to 2023: Schedule varies week to week, Same schedule every week, Choose own hours. The following amenities are derived from the Survey of Working Conditions and Psychosocial Risks (CT-RPS) from 2016 and 2019: Schedule – more visibility/notice, Schedule – more stability. Amenities derived from the CT-RPS are measured as standardized values of principal components. Amenities are observed at the firm-occupation-location-year level, where firm is the SIREN (enterprise) code, occupation is the 2-digit French PCS-ESE code, and location is the département.

drives identification in the model.

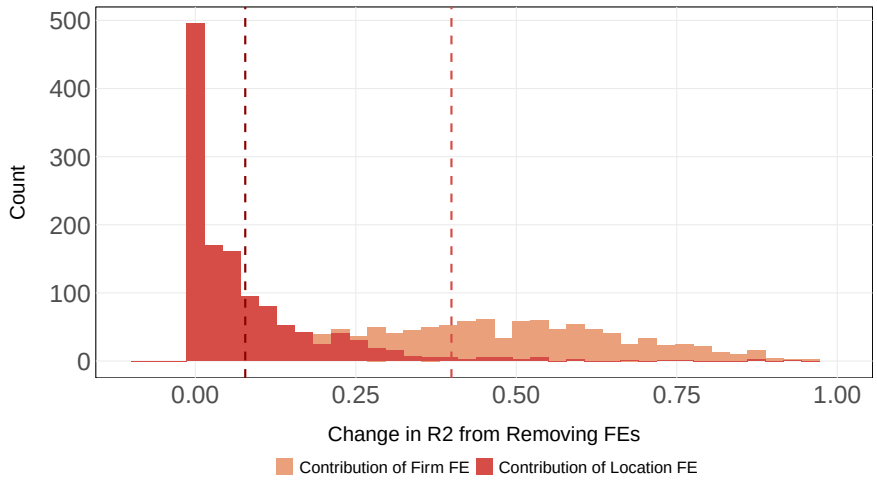
In this exercise, I build evidence that there is substantial variation in workplace amenities across firms within collective bargaining units. To do this, within a collective bargaining unit o, I regress the value of a given amenity a in firm j in location l on firm and location fixed effects:

$$a_{ojl} = \gamma_j + \gamma_l + \varepsilon_{ojl}, \tag{A4}$$

and repeat the regression twice: first dropping firm fixed effects γ_j and next dropping location fixed effects γ_l. I estimate the regression separately for each collective bargaining unit o. Regressions are weighted using the sample weights associated with each observation in the EEC. Then, for each CBA and each amenity, I calculate the change in R-squared associated with dropping firm or location fixed effects relative to the benchmark of the two-way fixed effect regression. The change in R-squared from dropping location fixed effects represents the marginal contribution of the location fixed effects; similarly, the change in R-squared from dropping firm fixed effects represents the

marginal contribution of firms. Figure B.2 shows the distribution of changes in R-squared values

FIGURE B.2. Firms Explain Variation in Amenities Within Collective Bargaining Units



Notes: This figure shows the marginal contributions of firm and location fixed effects to the R-squared values from regressions of the form shown in Equation A4. The marginal contributions are calculated by comparing the model with both types of one-way fixed effects to the model with jsut one dimension of fixed effects. The amenity measures are derived from the Continuous Labor Force Survey (EEC) from 2016 to 2023. Amenities are observed at the firm-occupation-location-year level, where firm is the SIREN (enterprise) code, occupation is the 2-digit French PCS-ESE code, and location is the département.

associated with dropping firm fixed effects and location fixed effects from the regression in Equation A4. The figure illustrates that there is substantial variation within CBA units across firms–the firm fixed effect explains an average of about 40 percent of the remaining variation.

Additionally, it is worth noting that in France, collective bargaining occurs at relatively decentralized levels. CBAs do not cover entire sectors as they do in many other countries. Because of this level of decentralization, even if amenity levels are somewhat restricted within CBAs, workers still face substantial heterogeneity across their choice sets because different jobs may belong to different collective bargaining units.

B.2. Firms, rather than locations, explain variation in occupation-specific amenities

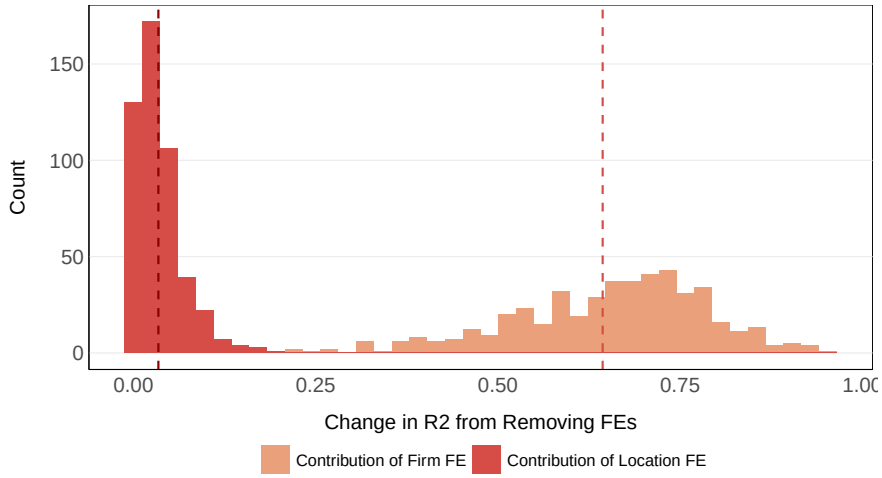
In this fact, I document that amenities are correlated within firms across space, conditional on occupation. To do this, I decompose the sources of variation in amenities across jobs, similar to Hazell et al. (2024)'s decomposition of variation in *wages*. This exercise leverages the idea that if amenities are correlated within firms, most of the variation in the job's characteristics will be explained by the firm rather than by the location. For each occupation o, we regress the value of a

given amenity a at firm j in location l on firm and location fixed effects:

$$a_{ojl} = \gamma_j + \gamma_l + \varepsilon_{ojl}, \tag{A5}$$

and repeat the regression twice: first dropping firm fixed effects γ_j and next dropping location fixed effects γ_l. We estimate the regression separately for each occupation o. Regressions are weighted using the sample weights associated with each observation in the EEC or RPS. Then, for each occupation and each amenity, I calculate the change in R-squared associated with dropping firm or location fixed effects relative to the benchmark of the two-way fixed effect regression. The change in R-squared from dropping location fixed effects represents the marginal contribution of the location fixed effects; similarly, the change in R-squared from dropping firm fixed effects represents the marginal contribution of firms.

FIGURE B.3. Occupation-Specific Amenities are Explained by Firms



Notes: This figure shows the marginal contributions of firm and location fixed effects to the R-squared values from regressions of the form shown in Equation A5. The marginal contributions are calculated by comparing the model with both types of one-way fixed effects to the model with jsut one dimension of fixed effects. The amenity measures are derived from the Continuous Labor Force Survey (EEC) from 2016 to 2023 and from the Survey of Working Conditions and Psychosocial Risks (CT-RPS) from 2016 and 2019. Amenities are observed at the firm-occupation-location-year level, where firm is the SIREN (enterprise) code, occupation is the 2-digit French PCS-ESE code, and location is the département.

Figure B.3 shows the distribution of changes in R-squared values associated with dropping firm fixed effects and location fixed effects from the regression in Equation A5. Each observation corresponds to an amenity-by-occupation specific regression. The figure shows that in every case the marginal contribution of location fixed effects in explaining amenity variation is minimal: the R-squared drops by less than 0.2 in all cases when location fixed effects are dropped, with a mean change in R-squared of 0.03. In contrast, firm fixed effects have a much larger marginal

contribution to variation in amenities within particular occupations. The mean of the distribution is 0.63, suggesting that, within occupations, more than half of the observed amenity variation occurs at the firm level with little variation across locations of the same firm.

B.3. Amenity choices exhibit spillovers within firms

Our next fact asks whether a particular job's amenities are influenced by the characteristics of other jobs at the same firm. Intuitively, if amenities are set at the firm-wide level, jobs should have characteristics that are correlated with those of *other* jobs *within* their own firms. In this fact, we document spillovers of amenity provision within a firm: conditional on one's own occupation, the occupations of one's colleagues influences the propensity of a given worker to adopt a particular amenity. Thus, workers' amenity takeup is influenced by colleagues' occupations through firm-wide policies.

To test this hypothesis, we focus on the case of remote work adoption. Remote work serves as an interesting case study because the large shock of the Covid-19 pandemic induced many firms to change their amenity policies. In labor markets more broadly, such large shocks are rare; amenities tend to evolve more incrementally over time. In this case, however, adoption of remote work rose sharply during the pandemic and stabilized at a persistently higher level in the following years. To develop an illustrative example of how firm-wide amenity setting would work in this context, consider the example of two accountants: one works at a technology firm where most of the other employees are software engineers; the other works at a wealth management firm where most of the other employees work in client relationship management. While the accountants' job tasks are similar — and are thus similarly predisposed to remote work — the tech accountant's colleagues in other positions are more likely to work remotely. Thus, we would expect that the tech accountant would be more likely to adopt remote work than the wealth management accountant if amenity policies spill over within the firm.

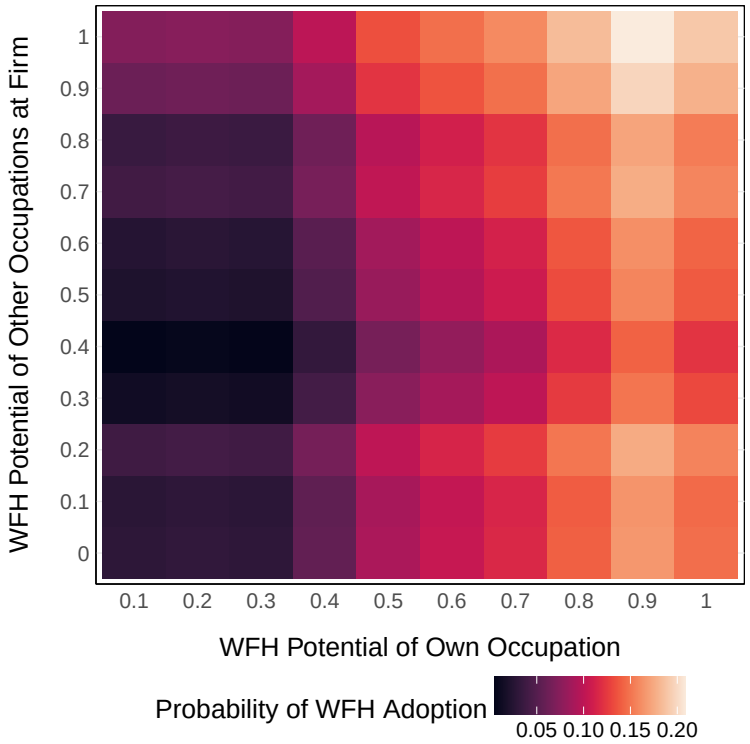
We exploit an additional advantage of the remote work context to develop this fact. A common issue in the study of amenities is that they are endogenously chosen, and may be correlated with firm-specific unobserved heterogeneity. Consequently, correlated amenity *adoption* within firms may reflect correlated unobservables rather than spillovers. To address this issue, we examine the effect of *potential* to work remotely of different occupations on *actual* remote work adoption. The outcome of interest y_{ojlt} is the share of workers in job oj (occupation o at firm j) who work from home for at least 50 percent of their working time. We regress y_{ojlt} on the probability that occupation o can be done from home, $P^{wfh}_{j,o}$, as well as the probability that other occupations $o-$ at firm j can be done

from home, $P^{wfh}_{j,o-}$:

$$y_{ojlt} = \alpha_o P^{wfh}_{j,o} + \beta_o P^{wfh}_{j,o-} + X_{ojlt} + \gamma_t + \varepsilon_{ojlt}, \tag{A6}$$

where X_{ojlt} is a vector of controls including the share of workers in job oj who are female and the mean age of workers in oj and γ_t are year fixed effects. Intuitively, the regression tests the hypothesis that, conditional on a worker's occupational potential to work remotely, the potential of her coworkers to do so will affect her realized propensity to take up remote work.

FIGURE B.4. Amenities Exhibit Cross-Occupation Spillovers within Firms



Notes: The measure of actual remote work adoption is derived from the Continuous Labor Force Survey (EEC) from 2016 to 2023. Telework adoption rates are observed at the firm-occupation-location-year level, where firm is the SIREN (enterprise) code, occupation is the 2-digit French PCS-ESE code, and location is the département. An individual is coded as working from home if they respond that they work from home either 100 percent of their working time or greater than 50 percent of their working time. The measure of potential to work from home is derived from data provided by Dingel and Neiman (2020), which classifies whether occupations can be done from home based on their task content in U.S. O*NET data. Occupations in the O*NET SOC classification are crosswalked to 4-digit ISCO-08 codes in the EEC data. Teleworkability of 4-digit ISCO codes corresponding to more than one O*NET SOC code are weighted based on the population frequency of the occupation to compute a probability that the worker is able to work remotely. Regressions are weighted using the EEC sample weights. Standard errors are clustered at the SIREN (enterprise) level.

Figure B.4 plots a heatmap of the estimated probability of adopting > 50 percent remote work for jobs with different *ex ante* WFH propensity and different levels of *colleagues' ex ante* WFH propensity. The heatmap shows a clear gradient from left to right in the probability of actual remote

work adoption. This shows that workers in jobs with higher propensity to work from home based on the task content of their occupations are more likely to actually work from home. However, there is also a gradient in the point estimates from top to bottom: conditional on a worker's own potential to work from home, being at a firm where workers in *other* occupations are more able to work from home makes it more likely that the job itself has a high rate of remote work takeup. The association is particularly pronounced when the leave-out remote work potential of other occupations is especially high. Even workers with low potential to work remotely are more likely to do so if their colleagues have near-perfect potential to work from home. Conversely, workers with high potential to work remotely will be less likely to do so if their colleagues in other occupations have little potential to work from home. This result suggests that firms may choose to set firm-wide amenity policies based on the composition of their workforce. This suggests a consequence of firm-wide amenity setting: the composition of one's colleagues may impose an externality – positive or negative – on the amenities of a particular job, beyond what can be rationalized by own-occupation characteristics alone.

B.3.1. Revealed preferences for jobs reflect both pay and amenities

Next, I ask whether observed amenities can account for residual variation in revealed preferences for firms that cannot be explained by wage premia alone. To answer this question, I leverage the revealed-preference measure of Sorkin (2018), known as PageRank, to infer the utility values of different jobs from employment-to-employment transitions. PageRank scores are measured at the job level, where a job is defined as an (enterprise, 2-digit occupation, region) tuple. Following exercises developed by Sorkin (2018) and Lagos (2024), I examine the relationship between the revealed preference scores of different jobs and their wage premia and observed amenities. I regress the PageRank score y_{ojl} on a vector of indicators for each BLM wage class $\boldsymbol{D_j^{BLM}}$, a vector of amenity measures $\boldsymbol{a_{ojl}}$, and fixed effects for the occupation γ_o or market-occupation ρ_{mo}:

$$y_{ojl} = \boldsymbol{\alpha \cdot D_j^{BLM}} + \beta \cdot \boldsymbol{a_{ojl}} + \gamma_o + \rho_{mo} + \varepsilon_{ojl}. \tag{A7}$$

I estimate the regression on two separate samples – one for the EEC and one for the CT-RPS – because there is very little overlap in jobs covered by both samples, making it infeasible to estimate one regression with all observed amenities. I cluster standard errors at the enterprise level, allowing for correlations in unobservables within firms, and weight regressions by the number of workers in each job in the BTS-Postes.

Table B.1 shows regression estimates for the EEC sample. Column (1) shows coefficients from the regression of PageRank score on indicators for each of the BLM wage premium classes and

TABLE B.1. PageRank Contributions of Firm Pay Premia and Amenities

a. EEC Sample

	(1)	(2)	(3)
BLM Class 2	0.278*	0.319	0.371*
	(0.167)	(0.210)	(0.218)
BLM Class 3	-0.323*	-0.391*	-0.372
	(0.188)	(0.226)	(0.271)
BLM Class 4	-0.215	-0.265	-0.233
	(0.162)	(0.202)	(0.212)
BLM Class 5	0.249	0.291	0.398*
	(0.181)	(0.216)	(0.220)
BLM Class 6	0.576***	0.654***	0.679***
	(0.191)	(0.225)	(0.236)
BLM Class 7	0.728***	0.801***	0.804***
	(0.201)	(0.239)	(0.244)
BLM Class 8	0.162	0.161	0.260
	(0.164)	(0.201)	(0.282)
BLM Class 9	0.593***	0.610***	0.697***
	(0.198)	(0.235)	(0.248)
BLM Class 10	0.527**	0.483*	0.847***
	(0.216)	(0.266)	(0.319)
Female share		-0.016	-0.255**
		(0.064)	(0.105)
> 50% WFH		-0.138*	-0.238***
		(0.076)	(0.092)
Work evenings		0.069	-0.022
		(0.056)	(0.081)
Work nights		-0.393***	-0.233**
		(0.082)	(0.115)
Choose hours			0.300***
			(0.100)
Occ. × Year FE	X	X	X
N	6,175	6,175	6,175
R-squared	0.106	0.143	0.165

Notes: This table reports results of regressions of the form described in Equation A7. The outcome variable in each column is the PageRank measure of revealed preferences described in Sorkin (2018). PageRank scores are computed at the job level, where a job is defined as an (enterprise, 2-digit occupation, region) tuple, using job-to-job transitions in the 2013-2021 Panel-BTS. BLM classes, or job wage premia, are computed following the approach of Bonhomme et al. (2019). I use k-means clustering with ten groups, clustering on 20 quantiles of the job wage distribution, weighted equally. Data is derived from the 2021-2023 EEC. Standard errors are clustered at the enterprise level.

occupation-by-year fixed effects. Columns (2) and (3) show coefficients from regressions with two sets of amenities. Column 2 uses a smaller set of amenities that are observed in all years of the EEC. Column 3 adds additional amenities that are asked about beginning in 2021, including whether the worker can choose her own working hours. In each of the regressions, I use the EEC sample from 2021-2023 to allow for a meaningful comparison of explanatory power across specifications.

The estimates in column (1) confirm that better-paying jobs are more desirable to workers: the job's revealed-preference score is increasing in wage class, with four of the five highest wage classes offering statistically significantly greater average utility than the lowest-paying class. However, the relationship is not strictly monotonic: jobs in classes 6, 9, and 10 have similar average PageRank values, with the highest-paying class having slightly *lower* utility than classes 6 and 9. Additionally, jobs in classes 3 and 4 are estimated to be less desirable than the lowest-paying class even after conditioning for occupation. The imperfect gradient suggests that wages matter to workers, but do not explain all of the value of a job.

After adding amenities in columns (2) and (3), the wage-utility gradient becomes steeper, particularly for the very highest-paying jobs. This result is consistent with amenities serving as compensating differentials: after controlling for non-wage characteristics, the utility premium for high-paying firms becomes *more* pronounced. For example, the coefficient for the top wage class increases by over 0.3 standard deviations from column (1) to column (3). The inclusion of amenities also substantially improves the explanatory power of the regression. The R-squared increases from 0.106 in the baseline model to 0.165 in the full specification, a 56% increase. This suggests that observed job-level amenities account for a substantial portion of residual variation in job desirability not captured by wages alone, and supports the idea that non-wage job characteristics are central to workers' revealed preferences.

The signs of the amenity coefficients offer intuitive interpretations. The coefficient on night work is strongly negative in both regressions, reflecting a clear disamenity. In contrast, being able to choose one's hours is strongly positively associated with revealed utility, highlighting the value of schedule control. Column (3) shows that jobs with a high female share are significantly less desirable, even conditional on pay and occupation, reflecting a "gender utility gap". Notably, the negative coefficient on the female share only becomes significant when autonomy (i.e., the ability to choose one's own hours) is controlled for, suggesting that women tend to sort into more flexible but otherwise less desirable jobs. The coefficients on remote work are negative and significant, suggesting that remote work is a disamenity. While the result may be intuitively surprising, it is consistent with hedonic wage regressions that tend to show that jobs with remote work command a pay premium (Mas, 2025). The sign may also reflect negative selection into remote work based on unobserved productivity (Emanuel and Harrington, 2024).

Appendix C. Mathematical Appendix

C.1. Derivations for a stylized model of firms' wage and amenity setting (Section 3.1)

The key features of firmwide amenity setting can be captured by a static model of the firm's production decision. To build intuition, we first consider a simplified version of the model in which (1) the firm uses two occupational inputs, $o = L, H$, and (2) the firm chooses a wage w_{oj} and a single amenity a_{oj} for each occupation o at time t. For simplicity, I abstract from the choice of capital inputs and suppress time subscripts.[15] The production function is given by

$$\begin{aligned} Y_j &= A_j \cdot \left((z_{Hj} \cdot L_{Hj})^{\rho_j} + (z_{Lj} \cdot L_{Lj})^{\rho_j} \right)^{\frac{\alpha_{Lj}}{\rho_j}}, \qquad \text{where} \\ z_{oj} &= \exp(\delta_{oj} \cdot a_{oj}) \qquad \text{for } o = L, H, \text{ and} \\ A_{jt} &= A_{j,t-1} + \nu_{jt}, \end{aligned} \tag{A8}$$

where $\nu_{jt} \sim N(0, \sigma_A)$. In each period, the firm's objective is to choose wages w_{oj} and amenities a_{oj} for each occupation $o = L, H$. We can summarize the firm's problem as

$$\max_{\{w_{oj}, a_{oj}\}_{o=L,H}} \left\{ A_j \cdot \left((z_{Hj} \cdot L_{Hj})^{\rho_j} + (z_{Lj} \cdot L_{Lj})^{\rho_j} \right)^{\frac{\alpha_{Lj}}{\rho_j}} - w_{Hj} L_{Hj} - w_{Lj} L_{Lj} \right\}.$$

The first order conditions for wages w_{oj} are given by

$$\frac{\alpha_{Lj}}{\rho_j} \cdot Y_j \cdot L_j^{-1} \cdot \rho_j \left((z_{Hj} L_{Hj})^{\rho_j - 1} \cdot z_{Hj} \frac{\partial L_{Hj}}{\partial w_{oj}} + (z_{Hj} L_{Hj})^{\rho_j - 1} \cdot z_{Lj} \frac{\partial L_{Lj}}{\partial w_{oj}} \right) = L_{oj} + w_{Hj} \frac{\partial L_{Hj}}{\partial w_{oj}} + w_{Lj} \frac{\partial L_{Lj}}{\partial w_{oj}}.$$

Define the elasticity of labor supply to job oj with respect to the same job's wage w_{oj} as

$$\varepsilon_{oj}^{w_{oj}} \equiv \frac{\partial L_{oj}}{\partial w_{oj}} \frac{w_{oj}}{L_{oj}} \quad \text{for } o = L, H.$$

Similarly, define the cross-elasticity of labor supply to job oj with respect to the other job at the same firm $o'j$ (i.e., the elasticity of L_L with respect to w_H and the elasticity of L_H with respect to w_L) as

$$\varepsilon_{Hj}^{w_{Lj}} \equiv \frac{\partial L_{Hj}}{\partial w_{Lj}} \frac{w_{Lj}}{L_{Hj}} \quad \text{and}$$

[15] It is straightforward to show that if production is Cobb-Douglas in aggregate labor and capital, the choices of wages does not depend on the choice of capital except through Y_j and the choice of amenity does not depend on capital at all. Thus, abstracting from the capital input decision does not affect the equilibrium equations for wages and amenities.

$$\varepsilon_{Lj}^{w_{Hj}} \equiv \frac{\partial L_{Lj}}{\partial w_{Hj}} \frac{w_{Hj}}{L_{Lj}}.$$

Then we can simplify the first order condition using the elasticities:

$$L_{Hj}(1+\varepsilon_{Hj}^{w_H}) + w_{Lj}\frac{\varepsilon_{Lj}^{w_{Hj}} L_{Lj}}{w_{Hj}} = \alpha_{Lj}\cdot Y_j \cdot L_j^{-1}\left((z_{Hj}L_{Hj})^{\rho_j}\cdot\frac{\varepsilon_{Hj}^{w_{Hj}}}{w_{Hj}} + (z_{Lj}L_{Lj})^{\rho_j}\cdot\frac{\varepsilon_{Lj}^{w_{Hj}}}{w_{Hj}}\right)$$

$$w_{Hj^*} = \frac{\alpha_{Lj}\cdot Y_j\cdot L_j^{-1}\left((z_{Hj}L_{Hj})^{\rho_j}\cdot\varepsilon_{Hj}^{w_{Hj}} + (z_{Lj}L_{Lj})^{\rho_j}\cdot\varepsilon_{Lj}^{w_{Hj}}\right) - w_{Lj}\cdot L_{Lj}\cdot\varepsilon_{Lj}^{w_{Hj}}}{L_{Hj}(1+\varepsilon_{Hj}^{w_{Hj}})}.$$

Next, the first order conditions for amenities a_{oj} are given by

$$\frac{\alpha_{Lj}}{\rho_j}\cdot Y_j\cdot L_j^{-1}\cdot\rho_j\times\left[(z_{Hj}L_{Hj})^{\rho_j-1}\cdot\left(z_{Hj}\frac{\partial L_{Hj}}{\partial a_{Hj}} + \delta_{Hj}z_{Hj}L_{Hj}\right) + (z_{Lj}L_{Lj})^{\rho_j-1}\cdot z_{Lj}\frac{\partial L_{Lj}}{\partial a_{Hj}}\right]$$
$$= w_{Hj}\frac{\partial L_{Hj}}{\partial a_{Hj}} + w_{Lj}\frac{\partial L_{Lj}}{\partial a_{Hj}}$$

Again, using the formulas for the labor supply elasticities with respect to the wages and analogous labor supply elasticities with respect to the amenities, we can rewrite:

$$\alpha_{Lj}\cdot Y_j\cdot L_j^{-1}\times\left[(z_{Hj}L_{Hj})^{\rho_j}\cdot\left(\frac{\varepsilon_{Hj}^{a_{Hj}}}{a_{Hj}}+\delta_{Hj}\right) + (z_{Lj}L_{Lj})^{\rho_j}\cdot\frac{\varepsilon_{Lj}^{a_{Hj}}}{a_{Hj}}\right] = w_{Hj}\frac{\varepsilon_{Hj}^{a_{Hj}}L_{Hj}}{a_{Hj}} + w_{Lj}\frac{\varepsilon_{Lj}^{a_{Hj}}L_{Lj}}{a_{Hj}}$$

$$\sum_{o=L,H} w_{oj}L_{oj}\varepsilon_{oj}^{a_{Hj}} - \alpha_{Lj}\cdot Y_j\cdot L_j^{-1}\sum_{o=L,H}(z_{oj}L_{oj})^{\rho_j}\varepsilon_{oj}^{a_{Hj}} = \alpha_{Lj}\cdot Y_j\cdot L_j^{-1}(z_{Hj}L_{Hj})^{\rho_j}\cdot\delta_{Hj}\cdot a_{Hj}$$

$$\frac{\sum_{o=L,H} w_{oj}L_{oj}\varepsilon_{oj}^{a_{Hj}}}{\alpha_{Lj}\cdot Y_j\cdot L_j^{-1}(z_{Hj}L_{Hj})^{\rho_j}\cdot\delta_{Hj}} - \frac{\sum_{o=L,H}(z_{oj}L_{oj})^{\rho_j}\varepsilon_{oj}^{a_{Hj}}}{(z_{Hj}L_{Hj})^{\rho_j}\cdot\delta_{Hj}} = a_{Hj}^*.$$

In this paper, I consider an equilibrium in which each firm is strategically small in the labor market, and thus ignores the effect of changing the wage in occupation H on labor supply to occupation L and vice versa. Formally, the assumption is:

ASSUMPTION A1 (Zero cross-job elasticities). $\varepsilon_{Hj}^{w_{Lj}} = \varepsilon_{Lj}^{w_{Hj}} = 0.$

This assumption is consistent with a framework in which (i) the firm views itself as infinitesimal within the labor market, and thus ignores the effects of changing wages on the wages of other jobs within the market, and (ii) jobs in different occupations are not close substitutes, so workers in one occupation are unlikely to respond to wage changes at jobs in a different occupation. This amounts to assuming that firms are "strategically small." The assumption is consistent with those made in

closely related papers in the literature, e.g., Lamadon, Mogstad, and Setzler (2022) and Kroft, Luo, Mogstad, and Setzer (2025).[16] Furthermore, I empirically find that the market shares of individual jobs are very small, supporting the choice to abstract from strategic interactions.

Under Assumption A1, the equilibrium wage becomes

$$w_{Hj^*} = \frac{\varepsilon_{Hj}^{w_{Hj}}}{1+\varepsilon_{Hj}^{w_{Hj}}} \cdot \alpha_{Lj} \cdot Y_j \cdot L_j^{-1} \cdot z_{Hj}^{\rho_j} \cdot L_{Hj}^{\rho_j - 1}, \tag{A9}$$

or, in logs, using that $z_{Hj} = \exp(\delta_{Hj} a_{Hj})$:

$$\log w_{Hj^*} = \log\left(\frac{\varepsilon_{Hj}^{w_{Hj}}}{1+\varepsilon_{Hj}^{w_{Hj}}}\right) + \log \alpha_{Lj} + \log\left(\frac{Y_j}{L_j}\right) + \rho_j \delta_{Hj} a_{Hj} + (\rho_j - 1)\log L_{Hj}. \tag{A10}$$

We can next simplify the expression for the equilibrium choice of amenity, focusing without loss of generality on a_{Hj}. Omitting the cross terms from the expression and then plugging in the equilibrium wage equation, we have

$$\begin{aligned} a_{Hj}^* &= \frac{w_{Hj} L_{Hj} \varepsilon_{Hj}^{a_{Hj}}}{\alpha_{Lj} \cdot Y_j \cdot L_j^{-1} (z_{Hj} L_{Hj})^{\rho_j} \cdot \delta_{Hj}} - \frac{(z_{Hj} L_{Hj})^{\rho_j} \varepsilon_{Hj}^{a_{Hj}}}{(z_{Hj} L_{Hj})^{\rho_j} \cdot \delta_{Hj}} \\ &= \frac{\varepsilon_{Hj}^{a_{Hj}}}{\delta_{Hj}} \left(\frac{\varepsilon_{Hj}^{w_{Hj}}}{1+\varepsilon_{Hj}^{w_{Hj}}} - 1\right) \\ a_{Hj}^* &= -\frac{\varepsilon_{Hj}^{a_{Hj}}}{(1+\varepsilon_{Hj}^{w_{Hj}})\delta_{Hj}}. \end{aligned} \tag{A11}$$

C.1.1. When will wages be correlated across occupations within the firm?

PROPOSITION A1. *Assume $\varepsilon_{Hjt}^{w_L} = 0$. Then w_{Hjt} and w_{Ljt} will have positive covariance if and only if the following two conditions hold:*

a. $\alpha_{Lj} > \rho_j$, *and*
b. $(\alpha_{Lj} - \rho_j) \cdot A > (1 - \rho_j) \cdot B$, *where*

$$A = \frac{(z_{Hjt} L_{Hjt})^{\rho_j}}{(z_{Hjt} L_{Hjt})^{\rho_j} + (z_{Ljt} L_{Ljt})^{\rho_j}} \cdot \frac{\varepsilon_{Ljt}^{w_L} \cdot w_{Hjt}}{w_{Ljt}}, \quad \textit{and}$$

[16] Recent papers testing for strategic interactions in wage setting (Sharma (2023), Roussille and Scuderi (2024), and Derenoncourt and Weil (2025)) do not find evidence of such interactions, suggesting that a model of monopsonistic competition as described here better captures firm behavior.

$$B = \frac{d \log L_{Hjt}}{dL_{Ljt}} \cdot L_{Ljt} \cdot \frac{\varepsilon^{w_L}_{Ljt} \cdot w_{Hjt}}{w_{Ljt}}.$$

PROOF. Recall the log wage equation for occupation o:

$$\log w^*_{ojt} = \log\left(\frac{\varepsilon^w_{ojt}}{1+\varepsilon^w_{ojt}}\right) + \log \alpha_{Lj} + \log\left(\frac{Y_{jt}}{L_{jt}}\right) + \rho_j \delta_{oj} a_{ojt} + (\rho_j - 1)\log L_{ojt}.$$

To derive the relationship $\partial w_H / \partial w_L$, observe that

$$\frac{\partial w_H}{\partial w_L} = \frac{\partial w_H}{\partial L_L} \cdot \frac{\partial L_L}{\partial w_L} + \frac{\partial w_H}{\partial L_H} \cdot \frac{\partial L_H}{\partial w_L},$$

and use the identities

$$\frac{\partial w_{Hj}}{\partial L_{Ljt}} = \frac{\partial w_{Hjt}}{\partial \log w_{Hjt}} \cdot \frac{\partial \log w_{Hjt}}{\partial L_{Ljt}},$$
$$\frac{\partial w_{Lj}}{\partial L_{Ljt}} = \frac{\partial w_{Ljt}}{\partial \log w_{Ljt}} \cdot \frac{\partial \log w_{Ljt}}{\partial L_{Ljt}}.$$

Using the fact that $dx/d\log(x) = x$, we can rewrite these as

$$\frac{\partial w_{Hj}}{\partial L_{Ljt}} = w_{Hjt} \cdot \frac{\partial \log w_{Hjt}}{\partial L_{Ljt}},$$
$$\frac{\partial w_{Lj}}{\partial L_{Ljt}} = w_{Ljt} \cdot \frac{\partial \log w_{Ljt}}{\partial L_{Ljt}}.$$

Next, using the wage equation, we can compute $\frac{\partial \log w_{Hjt}}{\partial L_{Ljt}}$ and $\frac{\partial \log w_{Ljt}}{\partial L_{Ljt}}$:

$$\begin{aligned}
\frac{d \log w_{Hjt}}{dL_{Ljt}} &= \frac{d}{dL_{Ljt}}\left(\log A_{jt} + \left(\frac{\alpha_{Lj} - \rho_j}{\rho_j}\right)\log\left((z_{Hjt}L_{Hjt})^{\rho_j} + (z_{Ljt}L_{Ljt})^{\rho_j}\right)\right) + (\rho_j - 1)\frac{d \log L_{Hjt}}{dL_{Ljt}} \\
&= (\alpha_{Lj} - \rho_j)\frac{z^{\rho_j}_{Ljt} L^{\rho_j - 1}_{Ljt}}{(z_{Hjt}L_{Hjt})^{\rho_j} + (z_{Ljt}L_{Ljt})^{\rho_j}} + (\rho_j - 1)\underbrace{\frac{d \log L_{Hjt}}{dL_{Ljt}}}_{\equiv \theta_{HL}} \\
\frac{d \log w_{Hjt}}{dL_{Hjt}} &= \frac{d}{dL_{Hjt}}\left(\log A_{jt} + \left(\frac{\alpha_{Lj} - \rho_j}{\rho_j}\right)\log\left((z_{Hjt}L_{Hjt})^{\rho_j} + (z_{Ljt}L_{Ljt})^{\rho_j}\right)\right) + (\rho_j - 1)\frac{d \log L_{Hjt}}{dL_{Hjt}} \\
&= (\alpha_{Lj} - \rho_j)\frac{z^{\rho_j}_{Hjt} L^{\rho_j - 1}_{Hjt}}{(z_{Hjt}L_{Hjt})^{\rho_j} + (z_{Ljt}L_{Ljt})^{\rho_j}} + (\rho_j - 1)\frac{1}{L_{Hjt}}
\end{aligned}$$

Replacing $\frac{\partial L_{Ljt}}{\partial w_{Ljt}} = \varepsilon^{w_L}_{Ljt} \cdot \frac{L_{Ljt}}{w_{Ljt}}$ and $\frac{\partial L_{Hjt}}{\partial w_{Ljt}} = \varepsilon^{w_L}_{Hjt} \cdot \frac{L_{Hjt}}{w_{Ljt}}$ in the expression for $\frac{\partial w_{Hjt}}{\partial w_{Ljt}}$, we have

$$\frac{\partial w_{Hjt}}{\partial w_{Ljt}} = (\alpha_{Lj} - \rho_j) \underbrace{\left(\frac{z_{Ljt}^{\rho_j} L_{Ljt}^{\rho_j}}{(z_{Hjt}L_{Hjt})^{\rho_j} + (z_{Ljt}L_{Ljt})^{\rho_j}} \cdot \frac{\varepsilon^{w_L}_{Ljt} \cdot w_{Hjt}}{w_{Ljt}} + \frac{z_{Hjt}^{\rho_j} L_{Hjt}^{\rho_j}}{(z_{Hjt}L_{Hjt})^{\rho_j} + (z_{Ljt}L_{Ljt})^{\rho_j}} \cdot \varepsilon^{w_L}_{Hjt} \right)}_{\equiv A} +$$
$$(\rho_j - 1) \underbrace{\left(w_{Hjt} \cdot \frac{d \log L_{Hjt}}{dL_{Ljt}} \cdot \varepsilon^{w_L}_{Ljt} \cdot \frac{L_{Ljt}}{w_{Ljt}} \right)}_{\equiv B} + (\rho_j - 1) \underbrace{\left(w_{Hjt} \cdot \frac{1}{L_{Hjt}} \cdot \varepsilon^{w_L}_{Hjt} \cdot \frac{L_{Hjt}}{w_{Ljt}} \right)}_{\equiv C}$$

To sign this expression, we make three observations. First, $A > 0$ so long as $\varepsilon^{w_L}_{Hjt}$ – the elasticity of labor supply to job (H, j) with respect to the wage in job (L, j) – is relatively small. Recall that we have assumed $\varepsilon^{w_L}_{Hjt} = 0$. Intuitively, if the labor markets for L and H are distinct, the elasticity will be close to zero – the elasticity will only be large if there is such significant overlap that raising wages in occupation L "cannibalizes" labor supply to occupation H.

Second, $\frac{d \log L_{Hjt}}{dL_{Ljt}} > 0$ if H and L are complements; similarly, $\frac{d \log L_{Hjt}}{dL_{Ljt}} < 0$ if L and H are substitutes. Third, $\rho_j < 1$ if and only if H and L are complements; likewise, $\rho_j > 1$ if and only if H and L are substitutes. Combining the second and third observations, $(\rho_j - 1)B \leq 0$ always (with equality iff $\rho_j = 1$), since $w_{Hjt}, w_{Ljt}, \varepsilon^{w_L}_{Ljt}$, and L_{Ljt} are all greater than zero.

Now, using our assumption that $\varepsilon^{w_L}_{Hjt} = 0$, we can simplify:

$$\frac{\partial w_{Hjt}}{\partial w_{Ljt}} = (\alpha_{Lj} - \rho_j) \underbrace{\left(\frac{z_{Ljt}^{\rho_j} L_{Ljt}^{\rho_j}}{(z_{Hjt}L_{Hjt})^{\rho_j} + (z_{Ljt}L_{Ljt})^{\rho_j}} \cdot \frac{\varepsilon^{w_L}_{Ljt} \cdot w_{Hjt}}{w_{Ljt}} \right)}_{\equiv A} +$$
$$(\rho_j - 1) \underbrace{\left(w_{Hjt} \cdot \frac{d \log L_{Hjt}}{dL_{Ljt}} \cdot \varepsilon^{w_L}_{Ljt} \cdot \frac{L_{Ljt}}{w_{Ljt}} \right)}_{\equiv B}$$
$$= (\alpha_{Lj} - \rho_j) \frac{z_{Ljt}^{\rho_j} L_{Ljt}^{\rho_j}}{(z_{Hjt}L_{Hjt})^{\rho_j} + (z_{Ljt}L_{Ljt})^{\rho_j}} + (\rho_j - 1) \left(\frac{d \log L_{Hjt}}{dL_{Ljt}} \cdot L_{Ljt} \right)$$
$$= (\alpha_{Lj} - \rho_j) \frac{z_{Ljt}^{\rho_j} L_{Ljt}^{\rho_j}}{(z_{Hjt}L_{Hjt})^{\rho_j} + (z_{Ljt}L_{Ljt})^{\rho_j}} + (\rho_j - 1) \left(\frac{d \log L_{Hjt}}{dL_{Hjt}} \frac{dL_{Hjt}}{dL_{Ljt}} \cdot L_{Ljt} \right)$$
$$= (\alpha_{Lj} - \rho_j) \frac{z_{Ljt}^{\rho_j} L_{Ljt}^{\rho_j}}{(z_{Hjt}L_{Hjt})^{\rho_j} + (z_{Ljt}L_{Ljt})^{\rho_j}} + (\rho_j - 1) \left(\frac{dL_{Hjt}}{dL_{Ljt}} \cdot \frac{L_{Ljt}}{L_{Hjt}} \right).$$

It follows immediately that $\frac{\partial w_{Hjt}}{\partial w_{Ljt}} > 0$ if and only if the conditions in the Proposition hold. □

C.1.2. When will amenities be correlated across occupations within the firm?

The first order conditions for amenities in each occupation o do not depend on Y_{jt} or L_{jt} as the conditions for wages do. Instead, each optimal amenity choice depends only on the elasticities of labor supply with respect to wages and amenities (ε^w_{ojt} and ε^a_{ojt}, respectively) and the productivity penalty associated with providing the amenity, δ_{oj}.

PROOF. Take logs of a_{ojt} and assume that $\delta_{oj} < 0$. We have

$$\log a_{ojt} = \log \varepsilon^a_{ojt} - \log(1 + \varepsilon^w_{ojt}) - \log(-\delta_{ojt}).$$

Then the covariance of $\log a_{Hjt}$ and $\log a_{Ljt}$ is given by

$$\begin{aligned}
\operatorname{Cov}(\log a_{Hjt}, \log a_{Ljt}) = &\underbrace{\operatorname{Cov}(\log \varepsilon^a_{Hjt}, \log \varepsilon^a_{Ljt})}_{\substack{\text{covariance of}\\ \text{amenity elasticities}}} + \underbrace{\operatorname{Cov}(\log(1 + \varepsilon^w_{Hjt}), \log(1 + \varepsilon^w_{Ljt}))}_{\substack{\text{covariance of}\\ \text{wage elasticities}}} - \\
&\underbrace{\operatorname{Cov}(\log \varepsilon^a_{Hjt}, \log(1 + \varepsilon^w_{Ljt})) - \operatorname{Cov}(\log(1 + \varepsilon^w_{Hjt}), \log \varepsilon^a_{Ljt})}_{\text{cross-covariance of wage and amenity elasticities}} + \\
&\underbrace{\operatorname{Cov}(\log(1 + \varepsilon^w_{Hjt}), \log(-\delta_{Lj})) + \operatorname{Cov}(\log(1 + \varepsilon^w_{Hjt}), \log(-\delta_{Lj}))}_{\text{cross-covariance of wage elasticities and efficiency of amenity provision}} - \\
&\underbrace{\operatorname{Cov}(\log \varepsilon^a_{Hjt}, \log(-\delta_{Lj})) - \operatorname{Cov}(\log(-\delta_{Hj}, \log \varepsilon^a_{Ljt}))}_{\text{cross-covariance of amenity elasticities and efficiency of amenity provision}} + \\
&\underbrace{\operatorname{Cov}(\log(-\delta_{Hj}), \log(-\delta_{Lj}))}_{\text{covariance of efficiencies of amenity provision}}.
\end{aligned}$$

To derive a condition under which $\operatorname{Cov}(\log a_{Hjt}, \log a_{Ljt}) > 0$, we can rearrange this expression:

$$\begin{aligned}
&\underbrace{\operatorname{Cov}(\log \varepsilon^a_{Hjt}, \log \varepsilon^a_{Ljt})}_{\substack{\text{covariance of}\\ \text{amenity elasticities}}} + \underbrace{\operatorname{Cov}(\log(1 + \varepsilon^w_{Hjt}), \log(1 + \varepsilon^w_{Ljt}))}_{\substack{\text{covariance of}\\ \text{wage elasticities}}} + \underbrace{\operatorname{Cov}(\log(-\delta_{Hj}), \log(-\delta_{Lj}))}_{\text{covariance of efficiencies of amenity provision}} + \\
&\underbrace{\operatorname{Cov}(\log(1 + \varepsilon^w_{Hjt}), \log(-\delta_{Lj})) + \operatorname{Cov}(\log(1 + \varepsilon^w_{Hjt}), \log(-\delta_{Lj}))}_{\text{cross-covariance of wage elasticities and efficiency of amenity provision}} > \\
&\underbrace{\operatorname{Cov}(\log \varepsilon^a_{Hjt}, \log(1 + \varepsilon^w_{Ljt})) + \operatorname{Cov}(\log(1 + \varepsilon^w_{Hjt}), \log \varepsilon^a_{Ljt})}_{\text{cross-covariance of wage and amenity elasticities}} +
\end{aligned}$$

$$\underbrace{\text{Cov}(\log \varepsilon^a_{Hjt}, \log(-\delta_{Lj})) + \text{Cov}(\log(-\delta_{Hj}, \log \varepsilon^a_{Ljt}))}_{\text{cross-covariance of amenity elasticities and efficiency of amenity provision}} \quad .$$

□

To evaluate whether this condition is likely to hold, we can sign each of the constituent terms. Below I provide an intuitive discussion of each sign. Note that upon estimating the key parameters I can empirically evaluate each of the covariances.

- **Covariance of amenity elasticities** ≥ 0**.** While there is not a great reason for amenity elasticities to be positively correlated, it is unlikely that they will be systematically *negatively* correlated.
- **Covariance of wage elasticities** > 0**.** Wage elasticities are related to the labor market power of the firm. If the firm has similar market power in multiple occupation-specific labor markets, markdowns will be positively correlated – and, thus, so will elasticities.
- **Covariance of efficiencies of amenity provision** > 0**.** If the productive efficiency of the firm depends at least in part on firm-specific technologies, this parameter will likely be correlated within the firm. For example, if the firm uses software to manage asynchronous schedules across workers in one occupation, it can use the same software for other occupations.
- **Cross-covariance of wage elasticities and amenity provision efficiency** < 0**.** If firm size is correlated with markdowns, large firms will face smaller wage elasticities. Larger firms are also likely to be more efficient in providing amenities. Thus if firm size is an important determinant of both of these parameters, firms with lower wage elasticities will have higher efficiencies of amenity provision, leading these terms to be negatively correlated.
- **Cross-covariance of wage and amenity elasticities** < 0**.** If workers with greater sensitivity to wages have less sensitivity to amenities, this term will be negative. However, it is plausible that some groups of workers will have more or less elastic labor supply globally, making them more responsive to both wages *and* amenities. If this is true, this cross-covariance will be positive.
- **Cross-covariance of amenity elasticities and amenity efficiency** > 0**.** If workers who are more sensitive to amenities sort into firms that are more efficient in providing those amenities, this covariance will be positive.

Notice that this does not require L and H to have similar productivity penalties δ^i_{oj}; for instance, δ^i_{Lj} could be very large and δ^i_{Hj} could be very small. But firms with *relatively* small penalties to amenity provision for high-skill workers will also have relatively small penalties to amenity provision for low-skill workers. This implies that amenities within the firm are correlated because of persistent technological efficiencies in providing them at the firm level. This result arises even without explicitly requiring the firm to provide the same amenity to each worker or imposing economies of scale in amenity provision within the firm.

C.2. Standard errors of substitution and commuting disutility parameters

Because the parameters of the nested logit labor supply model are estimated sequentially, we can only obtain the estimates of parameters at lower levels of the nesting structure after normalizing them by parameter estimates at higher levels. As a result, we also need to adjust the standard errors of the estimates at each level to account for this rescaling. We do so using the Delta Method. In this section, we outline the mapping between the quantities that are estimated directly at each level and the parameters of the model, and then derive formulas for the adjusted standard errors using estimated quantities.

We estimate λ_m directly in the top level of the problem and do not need to adjust it. At subsequent levels we estimate the following parameters $\hat{\theta}_{mo}$ and $\hat{\theta}_{\gamma}$:

$$\begin{aligned} \hat{\theta}_{mo} = \frac{\lambda_{mo}}{\lambda_m} \qquad &\Leftrightarrow \qquad \hat{\lambda}_{mo} = \hat{\theta}_{mo} \cdot \hat{\lambda}_m \\ \hat{\theta}_{\gamma} = \frac{\gamma_{mo}}{\lambda_{mo}} \qquad &\Leftrightarrow \qquad \hat{\gamma}_{mo} = \hat{\theta}_{\gamma} \cdot \hat{\lambda}_{mo} \end{aligned}$$

These relationships show how to derive the estimates of the parameters of interest $\hat{\lambda}_{mo}$ and $\hat{\gamma}_{mo}$. Next, we can derive the variance formulas for each of these objects using the Delta method:

$$\begin{aligned} Var(\lambda_{mo}) &\approx \left(\frac{\partial \lambda_{mo}}{\partial \theta_{mo}}\right)^2 Var(\hat{\theta}_{mo}) + \left(\frac{\partial \lambda_{mo}}{\partial \lambda_m}\right)^2 Var(\hat{\lambda}_m) \\ &= \hat{\lambda}_m^2 Var(\hat{\theta}_{mo}) + \hat{\theta}_{mo}^2 Var(\hat{\lambda}_m); \\ Var(\gamma_{mol}) &\approx \left(\frac{\partial \gamma_{mo}}{\partial \theta_{\gamma}}\right)^2 Var(\hat{\gamma}_{mo}) + \left(\frac{\partial \gamma_{mo}}{\partial \theta_{mo}}\right)^2 Var(\hat{\theta}_{mo}) + \left(\frac{\partial \gamma_{mo}}{\partial \lambda_m}\right)^2 Var(\hat{\lambda}_m) \\ &= \gamma_{mo}^2 \left[\left(\frac{Var(\hat{\gamma}_{mo})}{\theta_{\gamma}}\right)^2 + \left(\frac{Var(\hat{\theta}_{mo})}{\theta_{mo}}\right)^2 + \left(\frac{Var(\hat{\lambda}_m)}{\lambda_m}\right)^2\right]. \end{aligned}$$

C.3. Job-Level Labor Supply Elasticities

We are interested in computing the **job-level elasticity of labor supply for a worker in demographic group** d using the coefficient estimates from our nested logit model. Because preferences can vary according to the worker demographics d, it is natural to compute these elasticities separately. Accordingly, when I refer to the *share* of a particular job j, we refer to the share of workers in demographic group d who choose job j, s^d_{jmo}. In what follows, I repress time subscripts t to simplify notation.

C.3.1. Case taking markets as given (two-tiered nesting structure)

C.3.2. Case with endogenous market choice (three-tiered nesting structure)

In this derivation, I allow the worker to choose the market at the top level of the nesting structure. Then s_{ijmo}, the probability that worker i chooses job j in market m and occupation o, is a product of (1) the probability that i chooses market m, s_{im}; (2) the probability that i chooses occupation o conditional on choosing market m, $s_{io|m}$; and the probability that i chooses job j conditional on choosing occupation o and market m, $s_{ij|mo}$:

$$s_{ijmo} = s_{im} \cdot s_{io|m} \cdot s_{ij|mo}.$$

Then the derivative of s_{ijmo} with respect to the mean utility δ_{ij} that i assigns to j is given by

$$\frac{\partial s_{ijmo}}{\partial \mu_{ij}} = \frac{\partial s_{im}}{\partial \mu_{ij}} \cdot s_{io|m} \cdot s_{ij|mo} + s_{im} \cdot \frac{\partial s_{io|m}}{\partial \mu_{ij}} \cdot s_{ij|mo} + s_{im} \cdot s_{io|m} \cdot \frac{\partial s_{ij|mo}}{\partial \mu_{ij}}.$$

Then we can compute the individual partial derivatives with respect to each component, beginning with the bottom nesting level:

$$\begin{aligned}
\frac{\partial s_{ij|mo}}{\partial \mu_{ij}} &= \frac{\partial}{\partial \mu_{ij}} \left(\frac{\exp[(\mu_{ij} + \gamma_{mo} \cdot z_{ij})/\lambda_{m,o}]}{\sum_{j' \in \mathcal{J}_{mo}} \exp[(\mu_{ij'} + \gamma_{mo} \cdot z_{ij'})/\lambda_{m,o}]} \right) \\
&= \frac{1}{\lambda_{mo}} s_{ij|mo}(1 - s_{ij|mo}); \\
\frac{\partial s_{io|m}}{\partial \mu_{ij}} &= \frac{\partial s_{io|m}}{\partial I_{imo}} \cdot \frac{\partial I_{imo}}{\partial \mu_{ij}} \\
&= \frac{1}{\lambda_m} s_{io|m}(1 - s_{io|m}) \cdot s_{ij|mo}; \\
\frac{\partial s_{im}}{\partial \mu_{ij}} &= \frac{\partial s_{im}}{\partial I_{im}} \cdot \frac{\partial I_{im}}{\partial I_{imo}} \cdot \frac{\partial I_{imo}}{\mu_{ij}} \\
&= s_{im}(1 - s_{im}) \cdot s_{io|m} \cdot s_{ij|mo}.
\end{aligned}$$

Plugging into the original expression, we have

$$\begin{aligned}
\frac{\partial s_{ijmo}}{\partial \mu_{ij}} =& \frac{\partial s_{im}}{\partial \mu_{ij}} \cdot s_{io|m} \cdot s_{ij|mo} + s_{im} \cdot \frac{\partial s_{io|m}}{\partial \mu_{ij}} \cdot s_{ij|mo} + s_{im} \cdot s_{io|m} \cdot \frac{\partial s_{ij|mo}}{\partial \mu_{ij}} \\
=& s_{im}(1 - s_{im}) \cdot s_{io|m}^2 \cdot s_{ij|mo}^2 + s_{im} \cdot \frac{1}{\lambda_m} s_{io|m}(1 - s_{io|m}) \cdot s_{ij|mo}^2 + s_{im} \cdot s_{io|m} \cdot \frac{1}{\lambda_{mo}} \cdot s_{ij|mo}(1 - s_{ij|mo}) \\
=& s_{ijmo} \left((1 - s_{im}) \cdot s_{io|m} \cdot s_{ij|mo} + \frac{1}{\lambda_m}(1 - s_{io|m}) \cdot s_{ij|mo} + \frac{1}{\lambda_{mo}}(1 - s_{ij|mo}) \right). \qquad \text{(A12)}
\end{aligned}$$

Next, the derivative of s_{ijmo} with respect to the utility of job $k \neq j$ in the same occupation o and market m is

$$\frac{\partial s_{ijmo}}{\partial \mu_{ik}} = \frac{\partial s_{im}}{\partial \mu_{ik}} \cdot s_{io|m} \cdot s_{ij|mo} + s_{im} \cdot \frac{\partial s_{io|m}}{\partial \mu_{ik}} \cdot s_{ij|mo} + s_{im} \cdot s_{io|m} \cdot \frac{\partial s_{ij|mo}}{\partial \mu_{ik}} \cdot s_{ij|mo}.$$

And again we can take the partial derivatives of each constituent part:

$$\begin{aligned}
\frac{\partial s_{ij|mo}}{\partial \mu_{ik}} &= \frac{\partial}{\partial \mu_{ik}} \left(\frac{\exp[(\mu_{ij} + \gamma_{mo} \cdot z_{ij})/\lambda_{m,o}]}{\sum_{j' \in \mathcal{J}_{mo}} \exp[(\delta_{ij'} + \gamma_{mo} \cdot z_{ij'})/\lambda_{m,o}]} \right) \\
&= -s_{ij|mo} \frac{1}{\lambda_{mo}} s_{ik|mo}; \\
\frac{\partial s_{io|m}}{\partial \mu_{ik}} &= \frac{\partial}{\partial I_{imo}} \frac{\partial I_{imo}}{\partial \mu_{ik}} \\
&= \frac{1}{\lambda_m} s_{io|m}(1 - s_{io|m}) \cdot s_{ik|mo}; \\
\frac{\partial s_{im}}{\partial \mu_{ik}} &= \frac{\partial s_{im}}{\partial I_{im}} \frac{\partial I_{im}}{\partial I_{imo}} \frac{\partial I_{imo}}{\partial \mu_{ik}} \\
&= s_{im}(1 - s_{im}) \cdot s_{io|m} \cdot s_{ik|mo}.
\end{aligned}$$

Plugging back into the original expression, we have

$$\begin{aligned}
\frac{\partial s_{ijmo}}{\partial \mu_{ik}} =& s_{im}(1 - s_{im}) \cdot s_{io|m}^2 \cdot s_{ik|mo} \cdot s_{ij|mo} + \\
& s_{im} \cdot \frac{1}{\lambda_m} s_{io|m}(1 - s_{io|m}) \cdot s_{ik|mo} \cdot s_{ij|mo} - \\
& s_{im} \cdot s_{io|m} \cdot s_{ij|mo} \frac{1}{\lambda_{mo}} s_{ik|mo} \\
=& s_{im} \cdot s_{io|m} \cdot s_{ij|mo} \cdot s_{ik|mo} \left((1 - s_{im}) \cdot s_{io|m} + \frac{1}{\lambda_m}(1 - s_{io|m}) - \frac{1}{\lambda_{mo}} \right) \\
=& s_{ijmo} \cdot s_{ik|mo} \left((1 - s_{im}) \cdot s_{io|m} + \frac{1}{\lambda_m}(1 - s_{io|m}) - \frac{1}{\lambda_{mo}} \right).
\end{aligned} \tag{A13}$$

The slopes of the expected shares are the averages of the slopes across workers in demographic group d:

$$\frac{\partial s_{jmo}^d}{\partial \mu_k^d} = \frac{1}{N_d} \sum_{i=1}^{N_d} \frac{\partial s_{ijmo}}{\partial \mu_{ik}}$$

The general formula for the elasticity of the share of job j with respect to the wage at job k (whether $k = j$ or $k \neq j$) for workers in demographic group d is:

$$\frac{\partial \log s_{djmo}}{\partial \log w_{kmo}} = \frac{\alpha_d}{s^d_{jmo}} \frac{\partial s^d_{jmo}}{\partial \mu^d_k}$$

Then we can write the elasticity of labor supply to job j of a worker in demographic group d with respect to the wage at the same job j:

$$\varepsilon^w_{dj} = \frac{\alpha_d}{s^d_{jmo} N_d} \cdot \sum_{i=1}^{N_d} s_{ijmo} \left((1 - s_{im}) \cdot s_{io|m} \cdot s_{ij|mo} + \frac{1}{\lambda_m} (1 - s_{io|m}) \cdot s_{ij|mo} + \frac{1}{\lambda_{mo}} (1 - s_{ij|mo}) \right), \tag{A14}$$

and with respect to the wage at $k \neq j$:

$$\varepsilon_{jd}(w_{k \neq j}) = \frac{\alpha_d}{s^d_{jmo} N_d} \cdot \sum_{i=1}^{N_d} s_{ijmo} \cdot s_{ik|mo} \left((1 - s_{im}) \cdot s_{io|m} + \frac{1}{\lambda_m} (1 - s_{io|m}) - \frac{1}{\lambda_{mo}} \right). \tag{A15}$$

Analogously, to derive the elasticities of labor supply for workers in demographic group d to job j with respect to the amenity a_k at job k, we can write the general formula for the elasticity:

$$\frac{\partial \log s^d_{jmo}}{\partial \log w_{jmo}} = \frac{\beta^a_d}{s^d_{jmo}} \frac{\partial s^d_{jmo}}{\partial \mu^d_k}$$

Then we can write the elasticity of labor supply to job j of workers in demographic group d with respect to the amenity at the same job j:

$$\varepsilon^a_{dj} = \frac{\beta_d}{s^d_{jmo} N_d} \cdot \sum_{i=1}^{N_d} s_{ijmo} \left((1 - s_{im}) \cdot s_{io|m} \cdot s_{ij|mo} + \frac{1}{\lambda_m} (1 - s_{io|m}) \cdot s_{ij|mo} + \frac{1}{\lambda_{mo}} (1 - s_{ij|mo}) \right), \tag{A16}$$

and with respect to the amenity at $k \neq j$:

$$\varepsilon_{jd}(w_{k \neq j}) = \frac{\beta_d}{s^d_{jmo} N_d} \cdot \sum_{i=1}^{N_d} s_{ijmo} \cdot s_{ik|mo} \left((1 - s_{im}) \cdot s_{io|m} + \frac{1}{\lambda_m} (1 - s_{io|m}) - \frac{1}{\lambda_{mo}} \right). \tag{A17}$$

C.4. Equilibrium Conditions for Wages and Amenities

The firm's production function is Cobb-Douglas in capital K and occupation-specific labor inputs L_o. Firms j are allowed to vary in the set of occupations $\mathcal{O}_j$ that they take as inputs, as well as in thier output elasticities with respect to each input, α_{Kj}, α_{oj} for $o \in \mathcal{O}_j$. Then output Y_{jt} is given by:

$$Y_{jt} = A_{jt} \cdot K_{jt}^{\alpha_{Kj}} \cdot \prod_{o \in \mathcal{O}_j} \left(z_{ojt} L_{ojt}\right)^{\alpha_{oj}}, \tag{A18}$$

where z_{ojt} is a productivity penalty that depends on the provision of amenities and L_{ojt} is an aggregate of the efficiency units of labor in occupation o supplied by different worker productivity types X:

$$\begin{aligned} z_{ojt} &= \exp(\delta_{oj} a_{jt}) \\ L_{ojt} &= \int_X \theta_{ojX} L_{ojt}(X) dX. \end{aligned}$$

The firm's problem is to choose capital K_{jt}, a wage schedule $\{w_{ojt}(X)\}_{o,X}$ for occupations o and productivity types X, and an amenity a_{jt} to solve

$$\max_{K_{jt}, \{w_{ojt}(X)\}_{o,X}, a_{jt}} \left\{ A_{jt} K_{jt}^{\alpha_{Kj}} \prod_{o \in \mathcal{O}_j} \left(z_{ojt} L_{ojt}\right)^{\alpha_{oj}} - r_{jt} K_{jt} - \sum_{o \in \mathcal{O}_j} \int_X w_{ojt}(X) L_{ojt}(X) dX \right\}.$$

Evaluating the first order conditions gives us the equations for capital K_{jt}, the wage schedule $w_{ojt}(X)$, and the firm-level choice of amenity a_{jt}. We begin with the first order condition for capital, where we note that the firm takes the rental rate r_{jt} as fixed in period t:

$$\begin{aligned} \alpha_{Kj} K_{jt}^{-1} Y_{jt} - r_{jt} &= 0 \\ K_{jt}^{*} &= \frac{\alpha_{Kj} Y_{jt}}{r_{jt}}. \end{aligned}$$

Before proceeding, we make two assumptions on within-firm spillovers of changes to the wage for a particular occupation-type combination, $w_{ojt}(X)$, (1) to labor supplied by workers of other types X' and (2) to labor supplied by type X to other occupations $o' \neq o$ within the same firm j. We formalize these assumptions below before deriving the first order conditions for the wage schedule.

ASSUMPTION A2. *Let $L_{ojt}(X)$ be the units of labor supplied by workers of type X to occupation o at firm j and time t. We impose the following restrictions on the response of labor supply in different occupations $o' \neq o$ to changes in the*

a. $\frac{\partial L_{o'jt}(X)}{\partial w_{ojt}(X')} = 0$ *for* $X' \neq X$ *and for* $o' = o, o' \neq o$. *The labor supply of type X does not respond to changes in a wage for type* $X' \neq X$ *in any job at the same firm j.*

b. $\frac{\partial L_{o'jt}(X)}{\partial w_{ojt}(X)} = 0$ *for* $o' \neq o$. *The labor supply of type X to occupation o at firm j does not respond to changes in the wages offered to type X in other occupations* $o' \neq o$ *at the same firm j.*

Notice that we do *not* require an assumption on the response of labor supply in jobs at firm j to wages offered at any firm $j' \neq j$. The restrictions imposed by Assumption A2 simply require that there are no spillovers to the labor supply of other occupations and other types within the same firm.

We can now proceed to derive the first order conditions for the wages $w_{ojt}(X)$ – one for each (o, X) observed in the data:

$$\begin{aligned}
\sum_{o' \in \mathcal{O}_j} \alpha_{o'j} Y_{jt} L_{ojt}^{-1} \frac{\partial L_{o'jt}}{\partial w_{ojt}(X)} &= L_{ojt}(X) + \sum_{o' \in \mathcal{O}_j} \int_{X'} \left[w_{ojt}(X') \frac{\partial L_{o'jt}(X')}{\partial w_{ojt}(X)} \right] dX' \\
\sum_{o' \in \mathcal{O}_j} \alpha_{o'j} Y_{jt} L_{o'jt}^{-1} \frac{\partial L_{o'jt}}{\partial w_{ojt}(X)} &= L_{ojt}(X) + w_{ojt}(X) \frac{\partial L_{ojt}(X)}{\partial w_{ojt}(X)} \\
\alpha_{oj} Y_{jt} L_{ojt}^{-1} \theta_{ojX} \frac{\partial L_{ojt}(X)}{\partial w_{ojt}(X)} &= L_{ojt}(X) + w_{ojt}(X) \frac{\partial L_{ojt}(X)}{\partial w_{ojt}(X)}
\end{aligned}$$

where the simplification from the first line to the second line follows from part (a) of Assumption A2; the simplification from the second line to the third line follows from part (b) of the same. Now, to proceed, we define the labor supply elasticity with respect to the wage $\varepsilon^w_{ojX} \equiv \frac{\partial L_{ojt}(X)}{\partial w_{ojt}(X)} \frac{w_{ojt}(X)}{L_{ojt}(X)}$. Then we can substitute into the equation to eliminate partial derivatives:

$$w^*_{ojt}(X) = \frac{\varepsilon^w_{ojX}}{1 + \varepsilon^w_{ojX}} \alpha_{oj} \cdot Y_{jt} \cdot L_{ojt}^{-1} \cdot \theta_{ojX}.$$

Finally, we can derive the first order condition for the single amenity a_{jt} that is made available to all workers:

$$Y_{jt} \sum_{o \in \mathcal{O}_j} \left[\alpha_{oj} (z_{ojt} L_{ojt})^{\alpha_{oj}-1} \left(\delta_{oj} z_{ojt} L_{ojt} + z_{ojt} \int_X \theta_{ojX} \frac{\partial L_{ojt}(X)}{\partial a_{jt}} dX \right) \right] = \sum_{o \in \mathcal{O}_j} \int_X \left[w_{ojt}(X) \frac{\partial L_{ojt}(X)}{\partial a_{jt}} \right] dX$$

Similarly to how we defined the wage elasticity, we can define the elasticity of worker type (X) to job (o, j) with respect to the amenity a_{jt} as $\varepsilon^a_{ojX} \equiv \frac{\partial L_{ojt}(X)}{\partial a_{jt}} \frac{a_{jt}}{L_{ojt}(X)}$. Then we can substitute to eliminate partial derivatives:

$$A_{jt} K_{jt}^{\alpha_{Kj}} \sum_{o \in \mathcal{O}_j} \left[\alpha_{oj} (z_{ojt} L_{ojt})^{\alpha_{oj}-1} \prod_{o' \neq o} (z_{o'jt} L_{o'jt})^{\alpha_{o'j}} \left(\delta_{oj} z_{ojt} L_{ojt} + z_{ojt} \int_X \theta_{ojX} \varepsilon^a_{ojX} \frac{L_{ojt}(X)}{a_{jt}} dX \right) \right]$$

$$= \sum_{o \in \mathcal{O}_j} \int_X \left[w_{ojt}(X) \varepsilon^a_{ojX} \frac{L_{ojt}(X)}{a_{jt}} \right] dX$$

$$A_{jt} K_{jt}^{\alpha_{Kj}} \sum_{o \in \mathcal{O}_j} \left[\alpha_{oj} z_{ojt}^{\alpha_{oj}} L_{ojt}^{\alpha_{oj}-1} \prod_{o' \neq o} (z_{o'jt} L_{o'jt})^{\alpha_{o'j}} \left(\delta_{oj} L_{ojt} + \int_X \theta_{ojX} \varepsilon^a_{ojX} \frac{L_{ojt}(X)}{a_{jt}} dX \right) \right]$$

$$= \sum_{o \in \mathcal{O}_j} \int_X \left[w_{ojt}(X) \varepsilon^a_{ojX} \frac{L_{ojt}(X)}{a_{jt}} \right] dX$$

$$a_{jt} A_{jt} K_{jt}^{\alpha_{Kj}} \sum_{o \in \mathcal{O}_j} \left[\alpha_{oj} \delta_{oj} \prod_{o' \in \mathcal{O}_j} (z_{o'jt} L_{o'jt})^{\alpha_{o'j}} \right] + A_{jt} K_{jt}^{\alpha_{Kj}} \sum_{o \in \mathcal{O}_j} \int_X \theta_{ojX} \varepsilon^a_{ojX} L_{ojt}(X) dX$$

$$= \sum_{o \in \mathcal{O}_j} \int_X \left[w_{ojt}(X) \varepsilon^a_{ojX} L_{ojt}(X) \right] dX$$

$$a_{jt} Y_{jt} \sum_{o \in \mathcal{O}_j} \alpha_{oj} \delta_{oj} + A_{jt} K_{jt}^{\alpha_{Kj}} \sum_{o \in \mathcal{O}_j} \int_X \theta_{ojX} \varepsilon^a_{ojX} L_{ojt}(X) dX$$

$$= \sum_{o \in \mathcal{O}_j} \int_X \left[\frac{\varepsilon^w_{ojX}}{1 + \varepsilon^w_{ojX}} \alpha_{oj} \cdot Y_{jt} \cdot L_{ojt}^{-1} \cdot \theta_{ojX} \varepsilon^a_{ojX} L_{ojt}(X) \right] dX$$

To further simplify, we invoke an additional assumption: the elasticities of labor supply to occupation o at firm j with respect to the wage w do not vary with type X. We formalize this assumption below.

ASSUMPTION A3. $\varepsilon^w_{ojX} = \varepsilon^w_{oj}$ *for all* X.

Using this assumption, we can simplify further:

$$a_{jt} Y_{jt} \sum_{o \in \mathcal{O}_j} \alpha_{oj} \delta_{oj} + Y_{jt} \sum_{o \in \mathcal{O}_j} \alpha_{ojt} L_{ojt}^{-1} \int_X \theta_{ojX} \varepsilon^a_{ojX} L_{ojt}(X) dX$$

$$= Y_{jt} \sum_{o \in \mathcal{O}_j} \left[\frac{\varepsilon^w_{oj}}{1 + \varepsilon^w_{oj}} \alpha_{oj} \cdot L_{ojt}^{-1} \int_X \theta_{ojX} \varepsilon^a_{ojX} L_{ojt}(X) \right] dX$$

$$a_{jt} \sum_{o \in \mathcal{O}_j} \alpha_{oj} \delta_{oj} + \sum_{o \in \mathcal{O}_j} \alpha_{ojt} L_{ojt}^{-1} \int_X \theta_{ojX} \varepsilon^a_{ojX} L_{ojt}(X) dX$$

$$= \sum_{o \in \mathcal{O}_j} \left[\frac{\varepsilon^w_{oj}}{1 + \varepsilon^w_{oj}} \alpha_{oj} \cdot L_{ojt}^{-1} \int_X \theta_{ojX} \varepsilon^a_{ojX} L_{ojt}(X) \right] dX.$$

Next, define the **efficiency-weighted elasticity of labor supply to occupation o at firm j** as

$$\tilde{\varepsilon}^a_{oj} \equiv L_{ojt}^{-1} \int_X \theta_{ojX} \varepsilon^a_{ojX} L_{ojt}(X) dX. \tag{A19}$$

Intuitively, this object represents the elasticity of efficiency units of labor supply to occupation o

at firm j with respect to the amenity, taking into account the potential heterogeneity in elasticities across productivity types X and their differential productivity in occupation o, θ_{ojX}. Now we can simplify further to obtain an expression for the equilibrium amenity:

$$a_{jt} \sum_{o \in \mathcal{O}_j} \alpha_{oj} \delta_{oj} + \sum_{o \in \mathcal{O}_j} \alpha_{ojt} \tilde{\varepsilon}^a_{oj} dX = \sum_{o \in \mathcal{O}_j} \frac{\varepsilon^w_{oj}}{1 + \varepsilon^w_{oj}} \alpha_{oj} \cdot \tilde{\varepsilon}^a_{oj}$$

$$a^{jt*} = -\frac{\sum_{o \in \mathcal{O}_j} \frac{1}{1+\varepsilon^w_{oj}} \alpha_{oj} \cdot \tilde{\varepsilon}^a_{oj}}{\sum_{o \in \mathcal{O}_j} \alpha_{oj} \delta_{oj}}.$$

C.4.1. Extension: allowing amenities to vary across occupations, $a = a_{ojt}$

We can also solve a different version of the firm's problem in which the firm can choose a different level of the amenity a for each occupation, i.e., $a = a_{ojt}$. The first order conditions for capital and labor are the same as in the problem with a single choice of amenity. We now need to derive first order conditions for the occupation-specific amenities a_{ojt}. To do so, we invoke an additional assumption, similar to that of Assumption A2(b). We outline this assumption below.

ASSUMPTION A4. $\frac{\partial L_{o'jt}(X)}{\partial a_{ojt}} = 0$ *for* $o' \neq o$. *The labor supply of type X to occupation o at firm j does not respond to changes in the amenities offered to other occupations* $o' \neq o$ *at the same firm j.*

Now we can derive the first order conditions, recalling the definitions of ε^a_{ojX} and $\tilde{\varepsilon}^a_{ojX}$. They are given by

$$Y_{jt} \alpha_{oj} (z_{ojt} L_{ojt})^{-1} \left(\delta_{oj} z_{ojt} L_{ojt} + z_{ojt} \int_X \theta_{ojX} \frac{\partial L_{ojt}(X)}{\partial a_{ojt}} dX \right) = \int_X \left[w_{ojt}(X) \frac{\partial L_{ojt}(X)}{\partial a_{ojt}} \right] dX$$

$$Y_{jt} \alpha_{oj} L^{-1}_{ojt} \left(\delta_{oj} L_{ojt} + \int_X \theta_{ojX} \varepsilon^a_{ojX} \frac{L_{ojt}(X)}{a_{ojt}} dX \right) = \int_X \left[w_{ojt}(X) \varepsilon^a_{ojX} \frac{L_{ojt}(X)}{a_{ojt}} \right] dX$$

$$a_{ojt} Y_{jt} \alpha_{oj} \delta_{oj} + Y_{jt} \alpha_{oj} L^{-1}_{ojt} \int_X \theta_{ojX} \varepsilon^a_{ojX} L_{ojt}(X) dX = \int_X \left[w_{ojt}(X) \varepsilon^a_{ojX} L_{ojt}(X) \right] dX$$

$$a_{ojt} Y_{jt} \alpha_{oj} \delta_{oj} + Y_{jt} \alpha_{oj} L^{-1}_{ojt} \int_X \theta_{ojX} \varepsilon^a_{ojX} L_{ojt}(X) dX = \int_X \left[\frac{\varepsilon^w_{ojX}}{1 + \varepsilon^w_{ojX}} \alpha_{oj} \cdot Y_{jt} \cdot L^{-1}_{ojt} \cdot \theta_{ojX} \varepsilon^a_{ojX} L_{ojt}(X) \right] dX$$

$$a_{ojt} \alpha_{oj} \delta_{oj} + \alpha_{oj} L^{-1}_{ojt} \int_X \theta_{ojX} \varepsilon^a_{ojX} L_{ojt}(X) dX = \frac{\varepsilon^w_{oj}}{1 + \varepsilon^w_{oj}} \alpha_{oj} \cdot L^{-1}_{ojt} \int_X \left[\theta_{ojX} \varepsilon^a_{ojX} L_{ojt}(X) \right] dX$$

$$a_{ojt} \alpha_{oj} \delta_{oj} + \alpha_{oj} \tilde{\varepsilon}^a_{oj} = \frac{\varepsilon^w_{oj}}{1 + \varepsilon^w_{oj}} \alpha_{oj} \cdot \tilde{\varepsilon}^a_{oj}$$

$$a^*_{ojt} = -\frac{\tilde{\varepsilon}^a_{oj}}{(1 + \varepsilon^w_{oj}) \delta_{oj}}.$$

C.4.2. Extension: imperfect substitution of inputs with CES production

The Cobb-Douglas formulation of the production function assumes constant output elasticities in each of the inputs. What if we assume instead that inputs are imperfectly substitutable? In this section, we re-derive the equilibrium equations of the labor demand model after modifying the production to be CES in its labor inputs. The top-line production function captures substitution between capital and labor and remains Cobb-Douglas, i.e.:

$$Y_{jt} = A_{jt} \cdot K_{jt}^{\alpha_{Kj}} \cdot L_{jt}^{\alpha_{Lj}};$$

whereas the total quantity of labor supplied to firm j, L_{jt}, is now defined as a CES aggregate over all the occupations that j takes as inputs:

$$L_{jt} = \left(\sum_{o \in \mathcal{O}_j} (z_{ojt} L_{ojt})^{\rho_j} \right)^{\frac{1}{\rho_j}}.$$

As before, we aggregate the labor supply in occupation o at firm j linearly over worker productivity types X, and define z_{ojt} as before:

$$\begin{aligned} z_{ojt} &= \exp(\delta_{oj} a_{jt}) \\ L_{ojt} &= \int_X \theta_{ojX} L_{ojt}(X) dX. \end{aligned}$$

Now the firm's problem is to maximize profity by choice of capital K_{jt}, wage schedule $\{w_{ojt}(X)\}_{o,X}$, and amenity a_{ojt}:

$$\max_{K_{jt}, \{w_{ojt}(X)\}_{o,X}, \{a_{ojt}\}_o} \left\{ A_{jt} \cdot K_{jt}^{\alpha_{Kj}} \cdot \left(\sum_{o \in \mathcal{O}_j} (z_{ojt} L_{ojt})^{\rho_j} \right)^{\alpha_{Lj}/\rho_j} - r_{jt} K_{jt} - \sum_{o \in \mathcal{O}_j} \int_X w_{ojt}(X) L_{ojt}(X) dX \right\}.$$

First, note that because the top level is still Cobb-Douglas, the first order condition for capital K_{jt} is the same as in the original Cobb-Douglas specification:

$$K^* = \frac{\alpha_{Kj} Y_{jt}}{r_{jt}}.$$

Next, we can take first order conditions for the wages $\{w_{ojt}(X)\}_{o,X}$:

$$\frac{\alpha_{Lj}}{\rho_j} Y_{jt} L_{jt}^{-\rho} \left(\rho_{jt} (z_{ojt} L_{ojt})^{\rho_j - 1} \left(z_{ojt} \theta_{ojX} \frac{\partial L_{ojt}(X)}{\partial w_{ojt}(X)} \right) \right) = L_{ojt}(X) + w_{ojt}(X) \frac{\partial L_{ojt}(X)}{\partial w_{ojt}(X)}$$

$$\alpha_{Lj} Y_{jt} L_{ojt}^{-1} \left(\frac{z_{ojt} L_{ojt}}{L_{jt}} \right)^{\rho_j} \left(\theta_{ojX} \frac{\varepsilon^w_{ojtX}}{w_{ojt}(X)} \right) = 1 + \varepsilon^w_{ojtX}$$

$$w^*_{ojt}(X) = \frac{\varepsilon^w_{ojt}}{1 + \varepsilon^w_{ojt}} \cdot \alpha_{Lj} \cdot Y_{jt} \cdot L_{ojt}^{-1} \cdot \left(\frac{z_{ojt} L_{ojt}}{L_{jt}} \right)^{\rho_j} \cdot \theta_{ojX}.$$

The second line follows from the definition of the labor supply elasticity ε^w_{ojtX}; the third line follows from the previously invoked assumption that $\varepsilon^w_{ojtX} = \varepsilon^w_{ojt}$ for all X as well as some arithmetic cancellations. Finally, we can take the first order condition with respect to the amenities a_{ojt}:

$$\frac{\alpha_{Lj}}{\rho_j} Y_{jt} L_{jt}^{-\rho} \rho_j (z_{ojt} L_{ojt})^{\rho_j - 1} \left(\delta_{oj} z_{ojt} L_{ojt} + z_{ojt} \int_X \theta_{ojX} \frac{\partial L_{ojt}(X)}{\partial a_{ojt}} dX \right) = \int_X w_{ojt}(X) \frac{\partial L_{ojt}(X)}{\partial a_{ojt}} dX$$

$$\alpha_{Lj} Y_{jt} L_{jt}^{-\rho} (z_{ojt} L_{ojt})^{\rho_j - 1} \left(\delta_{oj} z_{ojt} L_{ojt} + z_{ojt} \int_X \theta_{ojX} \frac{\varepsilon^a_{ojtX} L_{ojt}(X)}{a_{ojt}} dX \right) = \int_X w_{ojt}(X) \frac{\varepsilon^a_{ojtX} L_{ojt}(X)}{a_{ojt}} dX$$

$$a_{ojt} \cdot \delta_{oj} + L_{ojt}^{-1} \int_X \theta_{ojX} \varepsilon^a_{ojtX} L_{ojt}(X) dX = \frac{\varepsilon^w_{ojt}}{1 + \varepsilon^w_{ojt}} L_{ojt}^{-1} \int_X \theta_{ojX} \varepsilon^a_{ojtX} L_{ojt}(X) dX$$

Again using the definition of $\tilde{\varepsilon}^a_{ojt} \equiv L_{ojt}^{-1} \int_X \theta_{ojX} \varepsilon^a_{ojtX} L_{ojt}(X) dX$, we can substitute to obtain the optimal choice of amenity:

$$a^*_{ojt} = -\frac{\tilde{\varepsilon}^a_{oj}}{(1 + \varepsilon^w_{ojt}) \delta_{oj}}.$$

Note that this expression for the choice of amenity is identical to the one obtained from the problem with Cobb-Douglas production. The key difference is in the wage equation, because the marginal product of the worker depends on the total labor supply in occupation o as well as the labor aggregate L_{jt}. In contrast, in the Cobb-Douglas case, only L_{ojt} explicitly enters the wage equation.

Appendix D. Measurement Appendix

D.1. Variable Definitions in Administrative Balance Sheet and Income Data

I follow the conventions developed by De Ridder (2024) for the definitions of variables from FARE used in production function estimation, summarized here.

- **Revenue** is total sales revenue. In FARE this is variable REDI_310.
- **Value added** is the economic value added net of factor costs. In FARE this is variable VACF. This corresponds to Y_{jt} in my model.
- **Employment** is total full-time-equivalent employment. In FARE this is variable REDI_E200. However, because the production function developed in the model takes as inputs the labor supplied to individual jobs (occupations and locations) within a firm j, I use the administrative matched employer-employee data (BTS-Postes) to measure full-time effective employment in each job rather than the equivalent variable from FARE. This is computed manually from employee counts and reported full-time percentages in the BTS-Postes.
- **Wage bill** is the sum of spending on salaries and wages (REDI_216) and on social security payments (REDI_217). I manually compute the firm's expenditures on salaries and wages for each occupation using gross wages in the BTS-Postes. I use the latter measure of the wage bill in place of REDI_216 when estimating the production function.
- **Direct production inputs** are calculated as the sum of purchases of merchandise and raw materials, adjusted for fluctuations in inventory. Following the definition of De Ridder (2024), this is calculated as the sum of purchases of goods (REDI_R210), changes in merchandise inventories (REDI_R211), purchases of raw materials and other supplies (REDI_R212), changes in inventories of raw materials and supplies (REDI_R213), and other purchases and external charges (REDI_R214).
- **Operating profits** is revenue minus the sum of the wage bill, material input expenses, capital depreciation, provisions, and other charges. This is equivalent to revenue (REDI_R310) minus total operating expenses (REDI_R201) in FARE.
- **Capital stock** is the sum of fixed tangible assets (IMMO_CORP) and intangible assets (IMMO_INCORP). We test robustness to the inclusion of only tangible assets. This corresponds to K_{jt} in my model.

D.2. Measurement of Job Characteristics in Survey Data

D.2.1. Continuous Labor Force Survey (EEC)

The first set of large-scale survey data on job characteristics comes from the *Échantillon Emploi en Continu* (EEC), or, in English, the Continuous Labor Force Survey. The EEC is the French administration of the common EU Labour Force Survey that is administered annually by all EU countries. The survey provides a generalized yet reasonably detailed description of the French labor market, collecting information on labor force participation, unemployment, and job search, as well as more detailed information on individual jobs from those who are currently employed. In its questions for employed workers, the EEC gathers information on employment, earnings, hours, occupations, and a variety of additional characteristics of jobs. The sample is quite large: in our preferred subsample of employed individuals aged 25 to 55, the annual sample varies in size from about 75,000 to about 110,000 over the years 2016 to 2023.

The EEC data currently cannot be linked at the individual level to other administrative data sources. However, it contains the French establishment and enterprise codes (SIRET and SIREN, respectively), the occupation of the worker using the French PCS-ESE classification scheme, and the location (département) of the job. We use these three variables to link survey data to other administrative files at the level of the job, where a distinct job is defined at the level of the enterprise $\times$ occupation $\times$ location. To aggregate individual survey data to the level of a job, we weight individual responses by the individual weights provided in the EEC (variable EXTRI). Table D.1 below summarizes the variables used to measure amenities in the EEC sample.

D.2.2. Survey of Working Conditions and Psychosocial Risks (CT-RPS)

The second set of large-scale survey data comes from the Enquête Conditions de travail et risques psychosociaux (CT-RPS) or, in English, the Survey of Working Conditions and Psychosocial Risks. The survey is conducted every three years; 2019 is the most recent year of the survey that is currently available. In our analysis, we use results from the 2016 and 2019 surveys. The survey aims to understand workplace conditions and risks, and asks a detailed set of questions similar to the O*NET survey in the United States.

Because the CT-RPS contains many questions that are relevant to workplace amenities, we follow Goldin (2014) by reducing the dimension of the data using a principal components analysis (PCA). The goal of the PCA is to group highly correlated job characteristics into summary measures that correspond to amenity quality in key topic areas. Tables A.1 and A.2 summarize the most important questions – defined by the magnitude of their factor loadings – in each of the first eight principal components. After deriving factor loadings using the PCA, we standardize values of the

TABLE D.1. Amenities and Disamenities in the Continuous Labor Force Survey (EEC)

ccc

Amenity	Variable	Question
Same schedule every week	HORAIC_Y	In your main job, your hours are: more or less the same from one week to the next
Schedule varies week to week	HORAIC_Y	In your main job, your hours are: variable from one week to the next
Choose own hours	DECIDHS_Y	Possibility to chose one's working schedule (entirely or to a certain extent)
Contract specified in days	FORFAIT_Y	Duration of work is determined in number of days to work, "forfait en jours"
Work from home > 50%	MAISOC_Y	Worked at home during the last 4 weeks: 100% of one's hours or 50% or more of one's hours
Work evenings	SOIRC_Y	Worked in the evening at any time during the last 4 weeks
Work nights	NUITC_Y	Worked at night at any time during the last 4 weeks

principal components within the survey data set so that values of the reduced-dimension amenities can be interpreted in terms of standard deviations from the mean. Topic areas corresponding to each principal components are chosen by author interpretation. PCs marked with a * indicate that the loadings on the principal component indicate that it corresponds to a disamenity. If a PC corresponds to a disamenity, its sign is reversed before standardizing so that positive values correspond to better amenities for each PC.

D.3. Measurement of Firm Wage Premia

I use the methodology of Bonhomme et al. (2019), henceforth BLM, to classify jobs' wage premia. A *job* is defined as a distinct enterprise (SIREN) × 2-digit occupation × location (departement). Jobs

are classified based on quantiles of their earnings distributions in the 2018 edition of the BTS-Postes, the French matched employer-employee data set that covers the universe of workers and firms in France in a given year. To compute quantiles of the job wage distributions, I first annualize earnings for each worker-job observation using the full-time-equivalency variable in the BTS-Postes data. The FTE variable accounts for the start and end dates of each position so that workers who are not present in the job for the full year can be considered as part of the job-wide wage distribution. For example, if a worker is recorded as working in a position for 6 months, her earnings are multiplied by 2 to impute an annual salary for the position. I then take logs of earnings before classifying. Then, for each job, I compute 20 quantiles of the log earnings distribution and apply k-means clustering to group together firms with similar earnings distributions.

I partition the set of J jobs into K earnings classes by solving the k-means problem:

$$\min_{k(1),...,k(J),H_1,...,H_K} \sum_{j=1}^{J} n_j \int \left(\hat{F}_j(y) - H_{k(j)}(y)\right)^2 d\mu(y) \tag{A20}$$

where $\hat{F}_j(y)$ is the (discretized) empirical cdf of log earnings in job j, n_j is the number of workers in job j, $k(1), ..., k(J)$ is the partition of jobs into K classes, an $H_1, ..., H_K$ are cdfs of the class wage distributions. Following BLM, I partition jobs into $K = 10$ classes. I also check robustness of the exercise to partitioning firms into $K = 15$ classes and find that the share of variation in earnings does not change much with more classes. Consequently, I prefer to use $K = 10$ for its greater computational tractability.

D.4. Measurement of Revealed Preferences for Firms/Jobs

I use the methodology of Sorkin (2018) to estimate the revealed-preference value of jobs in our sample, which infers the average utility of each job from employment-to-employment (E:E) transitions in panel data. Conceptually, the approach relies on the idea that "good jobs hire from other good jobs and have few workers leave." The vector of job values can be derived from observed employment-to-employment transitions in matched employer-employee data. In our preferred specification, we define *jobs* at the level of the enterprise × two-digit occupation × location. An enterprise is a firm, corresponding to a unique SIREN code in the French administrative data. Occupations are defined using the French PCS-ESE classification scheme. Locations are defined at the level of the département (territory) in which the job is located. We estimate revealed preferences for jobs rather than for firms in our main specification to allow for heterogeneity in the quality of jobs within the same firm. In what follows, we simplify notation by denoting each job j rather than explicitly specifying the triple (firm, location, occupation).

Define $\tilde{V}_j^{EE}$ as the common value of job j derived from E:E transitions. Next, define M_{jk}^o as the number of workers observed choosing between j and k who choose j (M_{kj}^o is defined analogously and $N = M_{kj}^o + M_{jk}^o$). If workers have Type I Extreme Value preferences over jobs (as I assume in the exposition of the model), such that workers' utilities over jobs $j = 1, ..., J$ are defined as $u_{ij} = \tilde{V}_j^{EE} + \zeta_{ij}$, then the probability of workers choosing j over k is

$$P(j \succ k) = \frac{\exp(\tilde{V}_j^{EE})}{\exp(\tilde{V}_j^{EE}) + \exp(\tilde{V}_k^{EE})},$$

and the ratio of common values can be estimated from the empirical choice probabilities:

$$\frac{M_{kj}^o}{N}\frac{N}{M_{jk}^N} = \frac{M_{kj}^o}{M_{jk}^N} = \frac{\exp(\tilde{V}_k^{EE})}{\exp(\tilde{V}_j^{EE})}. \tag{A21}$$

Equation A21 can be rewritten for jobs where we observe flows in both directions (such that M_{jk}^o and M_{kj}^o are both nonzero):

$$M_{kj}^o \exp(\tilde{V}_j^{EE}) = M_{jk}^o \exp(\tilde{V}_k^{EE}).$$

this holds for all jobs $j \in \mathcal{J}$, so we can sum over them:

$$\underbrace{\sum_{j\in\mathcal{J}} M_{kj}^o}_{\text{\# entering } k}\ \underbrace{\exp(\tilde{V}_j^{EE})}_{\text{value}} = \underbrace{\sum_{j\in\mathcal{J}} M_{jk}^o}_{\text{\# exiting } k}\ \underbrace{\exp(\tilde{V}_k^{EE})}_{\text{value}}, \tag{A22}$$

then dividing by the sum on the RHS gives us

$$\frac{\overbrace{\sum_{j\in\mathcal{J}} M_{kj}^o \exp(\tilde{V}_j^{EE})}^{\text{value-weighted entries}}}{\underbrace{\sum_{j\in\mathcal{J}} M_{jk}^o}_{\text{\# exiting } k}} = \underbrace{\exp(\tilde{V}_k^{EE})}_{\text{value of job } k}. \tag{A23}$$

Define the $J \times J$ diagonal **exit matrix** $\boldsymbol{S^o}$:

$$\boldsymbol{S^o} = \begin{bmatrix} \sum_j M^o_{j1} & 0 & \cdots & 0 \\ 0 & \sum_j M^o_{j1} & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \sum_j M^o_{jJ} \end{bmatrix}$$

Intuitively, the jth diagonal element of the exit matrix sums the number of workers exiting job j. Next, define the **mover matrix** $\boldsymbol{M^0}$:

$$\boldsymbol{M^0} = \begin{bmatrix} M_{11} & M_{12} & \cdots & M_{1J} \\ M_{21} & M_{22} & \cdots & M_{2J} \\ \vdots & \vdots & \ddots & \vdots \\ M_{J1} & M_{J2} & \cdots & M_{JJ} \end{bmatrix}$$

The i, jth element of the mover matrix is the number of workers choosing between i and j who chooose i. The vector of job-level values $\exp(\tilde{V}^{EE})$ is the unique fixed point of the function $\boldsymbol{(S^o)^{-1}M^o} : |\mathbb{R}|^{\mathcal{J}} \to |\mathbb{R}|^{\mathcal{J}}$:

$$\boldsymbol{(S^o)^{-1}M^o} \exp(\tilde{V}^{EE}) = \exp(\tilde{V}^{EE}). \tag{A24}$$

Note that with the definitions of $\boldsymbol{S^o}$ and $\boldsymbol{M^0}$ provided here, equation A24 is simply a restatement of equation A23, stacked across all J jobs, in matrix notation. Sorkin (2018) also derives rescalings of the M_{ij} values to account for differences in job size.

D.5. Network Clustering of Labor Markets using Job-to-Job Transitions

To make the job choice model both conceptually realistic and computationally tractable, I partition the economy into a set of distinct labor markets. Workers in our model do not choose from the universe of jobs in the economy; having them do so would not only misrepresent the structure of actual search behavior but would also make estimation infeasible due to the sheer number of alternatives. A natural solution is to restrict the set of job options available to each worker by defining discrete labor markets. But how should these boundaries be drawn?

A common approach is to define labor markets using observed classifications, such as occupation codes or geographic regions. However, this method imposes strong assumptions about the underlying structure of job mobility and can be sensitive to arbitrary choices like the level of occupational aggregation (e.g., 2-digit vs. 3-digit codes) or the classification system used (e.g., the French

PCS-ESE system used in this paper vs. the standardized European ISCO codes). Moreover, the relevance of geographic or occupational boundaries may vary substantially across worker groups. As Nimczik (2023) notes, while geographic areas are fixed across individuals, workers can differ greatly in how they perceive or act on job mobility opportunities. In other words, the true boundaries of a labor market may differ across workers and across job types.

To address these concerns, I follow the empirical strategy of Nimczik (2023), who proposes a data-driven approach to identifying labor markets based on observed job mobility. The idea is to use panel data on job-to-job transitions to construct a network of jobs, where jobs serve as nodes and observed transitions between jobs define directed, weighted edges. Rather than relying on pre-defined groupings, this approach infers the latent labor market structure from the observed flow of workers. The network is then partitioned into distinct labor markets using a stochastic block model (SBM)—a standard approach in the network literature for detecting latent communities (Holland et al. 1983; Karrer and Newman 2011). In this framework, labor markets correspond to communities of jobs that are tightly linked through mobility patterns.

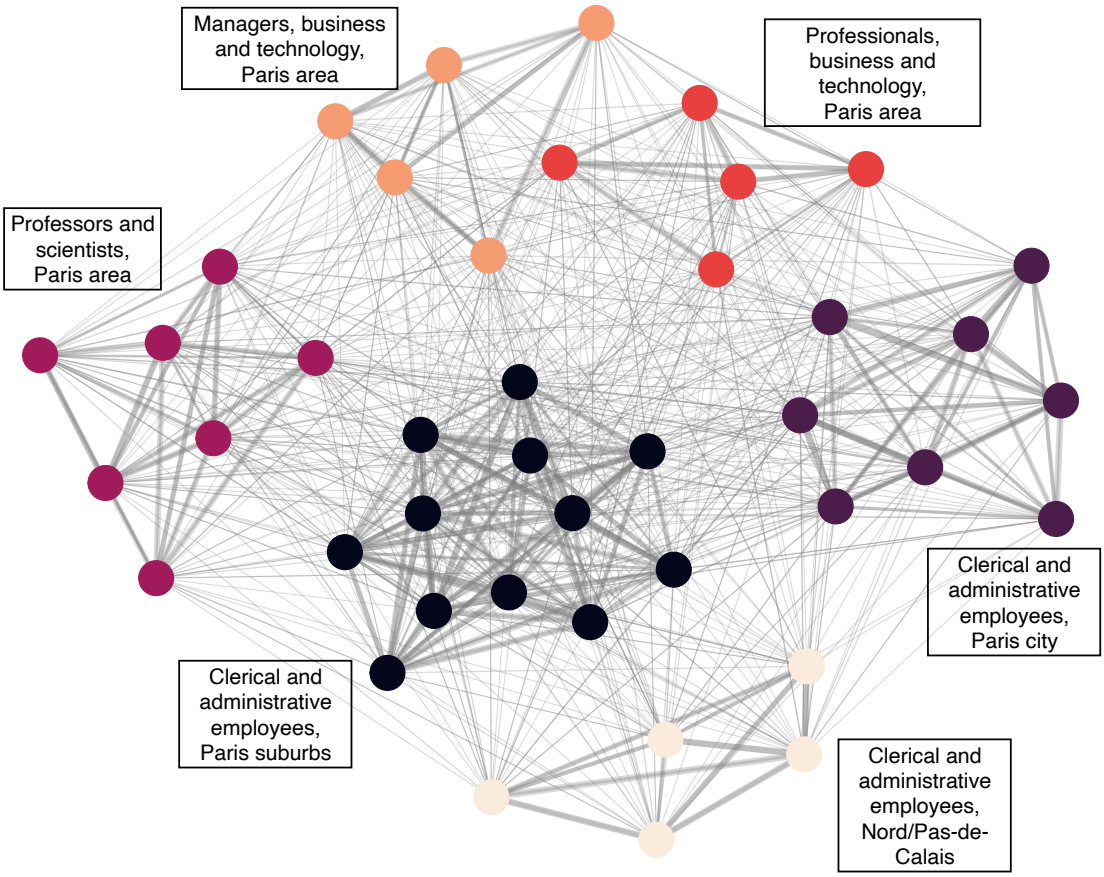
Importantly, this method has a structural interpretation. As Nimczik (2023) emphasizes, the network graph can be viewed as a reduced-form analogue to a gravity model of job transitions, where the likelihood of a worker moving between two jobs depends on the relative costs of switching. Intuitively, the cost of switching from one job to another is low if the two jobs are similar and high if they are different. This provides a microfoundation for the network structure that echoes work by Eaton and Kortum (2002) and Cortes and Gallipoli (2018). The resulting jobs clusters can be interpreted as latent labor markets. I refer the interested reader to Nimczik (2023) for greater detail on the structural economic interpretation of the network structure.

The set of job-to-job transitions is measured over the Panel-BTS file period of 2013-2021. Each **node** in the network is a single *job*: an enterprise-occupation-location tuple (j, o, l). Each **directed, weighted edge** in the network connecting two jobs (j, o, l) and (j', o', l') is number of times a worker is observed moving from (j, o, l) to (j', o', l'). I define job-to-job transitions in the 2013-2021 panel using the following restrictions. First, I include only transitions with no more than 30 days between the end of the first job and the start of the second job. Second, I require that workers hold full-time positions in each job and have tenure in each of at least one year. Third, I exclude transitions where either job is a temporary or apprenticeship position, though jobs can be either permanent or fixed-term contracts. Fourth, I require both jobs to be at firms with at least five employees. Finally, I drop transitions for workers born before 1945 or after 1996 (who are either beyond retirement age or under the age of 25 in the whole panel) and workers for whom I only ever observe a single employment spell.

To illustrate how the procedure partitions the economy into discrete labor markets, I visualize

the partition in Figure D.1 below. I generate the nodes and edges using simulated data for illustrative purposes only. The interpretations of the different clusters as labelled on the figure are drawn from actual labor markets identified in the data.

FIGURE D.1. Illustrative example of labor mobility network and market clusters



Notes: TK.

One might be concerned that the market partition may be sensitive to the definitions of jobs or to the inclusion of transitions between jobs within the same firm. To probe the robustness of the classification, I compute the network partition using eight specifications: defining jobs using two-, three-, and four-digit occupations and including or excluding within-firm job transitions. I can ask how similar the partitions are under each of these classification schemes using the normalized mutual information (NMI) measure for each pairwise comparison of candidate schemes (Danon et al., 2005). Intuitively, NMI measures the extent to which knowing the label of a node under one partition tells us what the label will be under the other partition. A score of 0 indicates that the partitions are completely mutually uninformative. A score of 1 indicates that knowing the label

under one partition determines the label under the other.

Table D.4 above reports the pairwise NMI scores for each combination of candidate partitions. By construction, comparisons of the same partition have NMI scores of 1. All other permutations report high NMI scores, with a minimum of 0.768 and the majority falling above 0.85. These results suggest that the classification of jobs is robust to the choice of more or less granular occupations. With this in mind, our preferred market classification is the one using 2-digit occupation and including within-firm transitions, because it includes the largest set of workers and jobs out of the available partitions.

TABLE D.2. Constituent Questions of Principal Components 1-4 of CT-RPS Data

Question Rank in PC	Loading	Question Text (English)
1. Culture – Respect		
1	0.24	I receive the respect that my work deserves
2	0.22	I am treated in an equitable way at work
3	0.21	I experience unpredictable or poorly prepared changes
4	0.21	My superior pays attention to what I say
5	0.21	I work under pressure
6	0.21	I have to do things I disapprove of
7	0.20	I have to hide my emotions
8	0.20	I can organize my work in the way that suits me best
9	0.20	Those who evaluate my work know it ell
10	0.19	I have the opportunity to develop my professional skills
2. Hazards*		
1	0.28	Carrying or moving heavy objects
2	0.28	Standing for long periods
3	0.27	Strenuous our tiring movements
4	0.24	Make long or frequent journeys on foot
5	0.23	Contact with dangerous products
6	0.22	works in the evening – often
7	0.22	schedule varies week to week
8	-0.22	schedule is the same week to week
9	0.17	Remain in an uncomfortable or tiring position for long periods
10	0.16	works at night – often
3. Stress/overwork*		
1	0.29	Work outside of anticipated hours
2	0.26	Bring work home
3	0.24	Contacted outside of working hours for work needs during the last twelve months
4	-0.22	Sufficient time to do my job properly
5	0.21	I am asked to do an excessive amount of work
6	0.21	Tension in relationships with the public
7	-0.20	I think about work even when I'm not there
8	-0.20	Breathe fumes or dust
9	0.20	Tension in relationships with subordinates
10	0.20	works in the evening – sometimes
4. Schedule – variability*		
1	-0.33	schedule is the same week to week
2	0.33	schedule varies week to week
3	-0.28	Knowledge of working hours a month in advance
4	0.23	Knowledge of working hours a week in advance
5	-0.22	I sometimes fear for my or others' safety
6	-0.22	I am asked to do an excessive amount of work
7	0.18	Sufficient time to do my job properly
8	-0.18	I have to do things I disapprove of
9	-0.18	I can organize my work in the way that suits me best
10	0.17	Sufficient and appropriate continuing training to do my job properly

Notes: This table summarizes the first 10 survey questions with the greatest magnitude of factor loadings in the first four principal components (PCs) of data from the 2016 Survey of Working Conditions and Psychosocial Risks (CT-RPS). PCs are ranked in order of the share of total variation in the survey data that they explain. Topic areas corresponding to each principal components are chosen by author interpretation. PCs marked with a * indicate that the loadings on the principal component indicate that it corresponds to a disamenity. If a PC corresponds to a disamenity, its sign is reversed before standardizing so that positive values correspond to better amenities for each PC.

TABLE D.3. Constituent Questions of Principal Components 5-8 of CT-RPS Data

Question Rank in PC	Loading	Question Text (English)
1. Schedule – Visibility/Notice		
1	0.35	Knowledge of working hours a month in advance
2	0.28	Receiving help from colleagues
3	0.27	Opportunity to collectively address questions of organization or operation
4	0.25	Organized meetings to collectively address questions
5	-0.24	Knowledge of working hours a week in advance
6	-0.22	I work alone
7	-0.22	Knowledge of working hours a day in advance
8	0.22	The possibility to cooperate to do my job correctly
9	0.19	Receiving help from superiors
10	0.19	schedule is the same week to week
6. Odd Hours*		
1	-0.42	works in the evening – sometimes
2	0.40	works in the evening – often
3	-0.37	works at night – sometimes
4	0.34	works at night – often
5	-0.24	Breathe fumes or dust
6	-0.24	Jolts or vibrations
7	-0.20	Possibility to temporarily interrupt work whenever desired
8	-0.19	Tension in relationships with subordinates
9	0.16	Knowledge of working hours a month in advance
10	-0.15	Contact with dangerous products
7. Flexibility		
1	0.32	Bring work home
2	-0.24	work affects health in a negative way
3	-0.23	Possibility to temporarily interrupt work whenever desired
4	-0.22	Need to follow strict quality control processes
5	0.22	Standing for long periods
6	-0.20	I am asked to do an excessive amount of work
7	0.20	Feeling of doing something useful to others
8	-0.19	I sometimes fear for my or others' safety
9	0.19	Knowledge of working hours a month in advance
10	-0.18	Feeling of being exploited
8. Culture – Relationships		
1	-0.39	Tension in relationships with colleagues
2	-0.33	Tension in relationships with subordinates
3	-0.32	Tension in relationships with superiors
4	-0.24	Sufficient time to do my job propeerly
5	0.23	I am asked to do an excessive amount of work
6	0.21	Receiving help from superiors
7	0.20	Receives strict operating instructions
8	0.18	Receiving help from colleagues
9	0.17	Knowledge of working hours a week in advance
10	-0.17	schedule varies week to week

Notes: This table summarizes the first 10 survey questions with the greatest magnitude of factor loadings in the fifth through eighth principal components (PCs) of data from the 2016 Survey of Working Conditions and Psychosocial Risks (CT-RPS). PCs are ranked in order of the share of total variation in the survey data that they explain. Topic areas corresponding to each principal components are chosen by author interpretation. PCs marked with a * indicate that the loadings on the principal component indicate that it corresponds to a disamenity. If a PC corresponds to a disamenity, its sign is reversed before standardizing so that positive values correspond to better amenities for each PC.

TABLE D.4. Pairwise normalized mutual information (NMI) by occupation and inclusion of within-firm job changes

Occupation detail		**2-digit**		**3-digit**		**4-digit**	
Incl. within-firm transitions?		**No**	**Yes**	**No**	**Yes**	**No**	**Yes**
2-digit	**No**	1.000	0.826	0.906	0.870	0.869	0.883
	Yes	0.826	1.000	0.792	0.877	0.768	0.834
3-digit	**No**	0.906	0.792	1.000	0.854	0.939	0.891
	Yes	0.870	0.877	0.854	1.000	0.827	0.898
4-digit	**No**	0.869	0.768	0.939	0.827	1.000	0.879
	Yes	0.883	0.834	0.891	0.898	0.879	1.000

Notes: TK.

FIGURE D.2. Sector Composition of Data-Driven Labor Markets

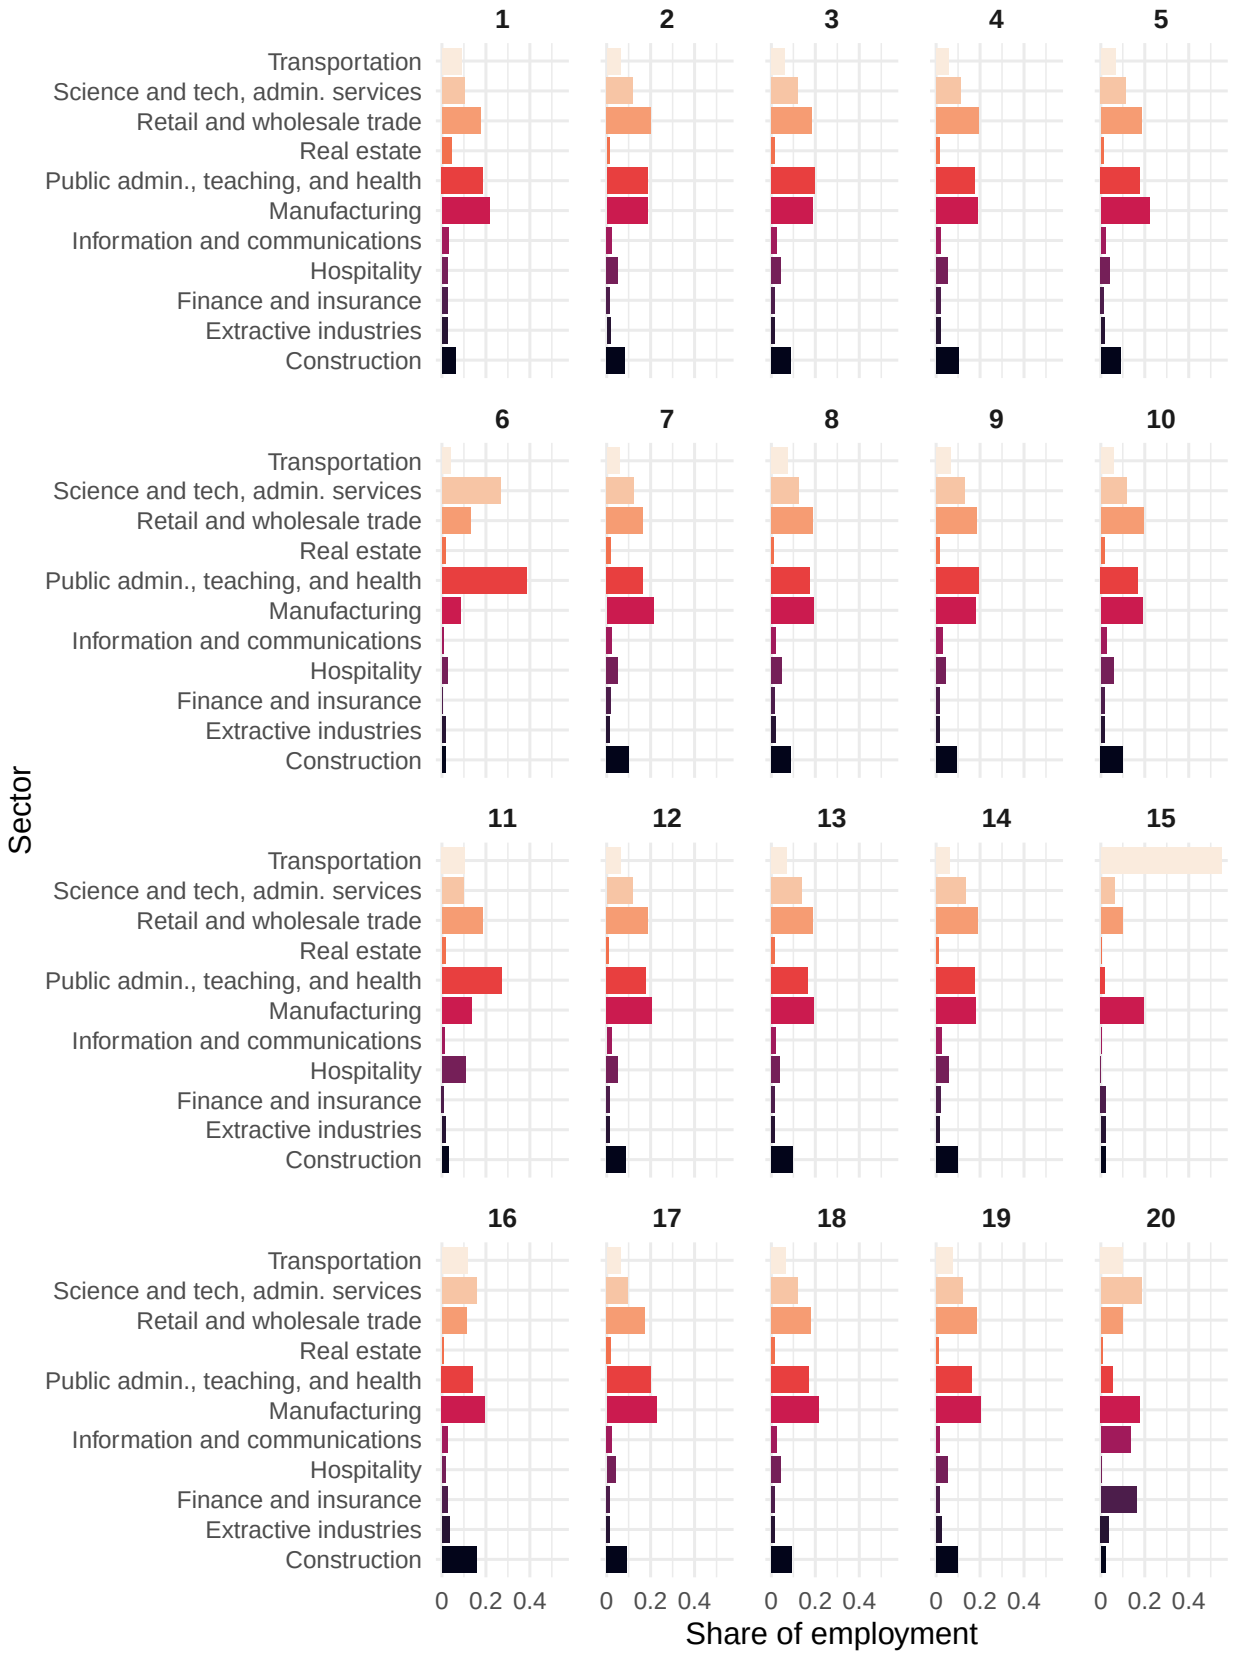


FIGURE D.3. 1-Digit Occupation Composition of Data-Driven Labor Markets

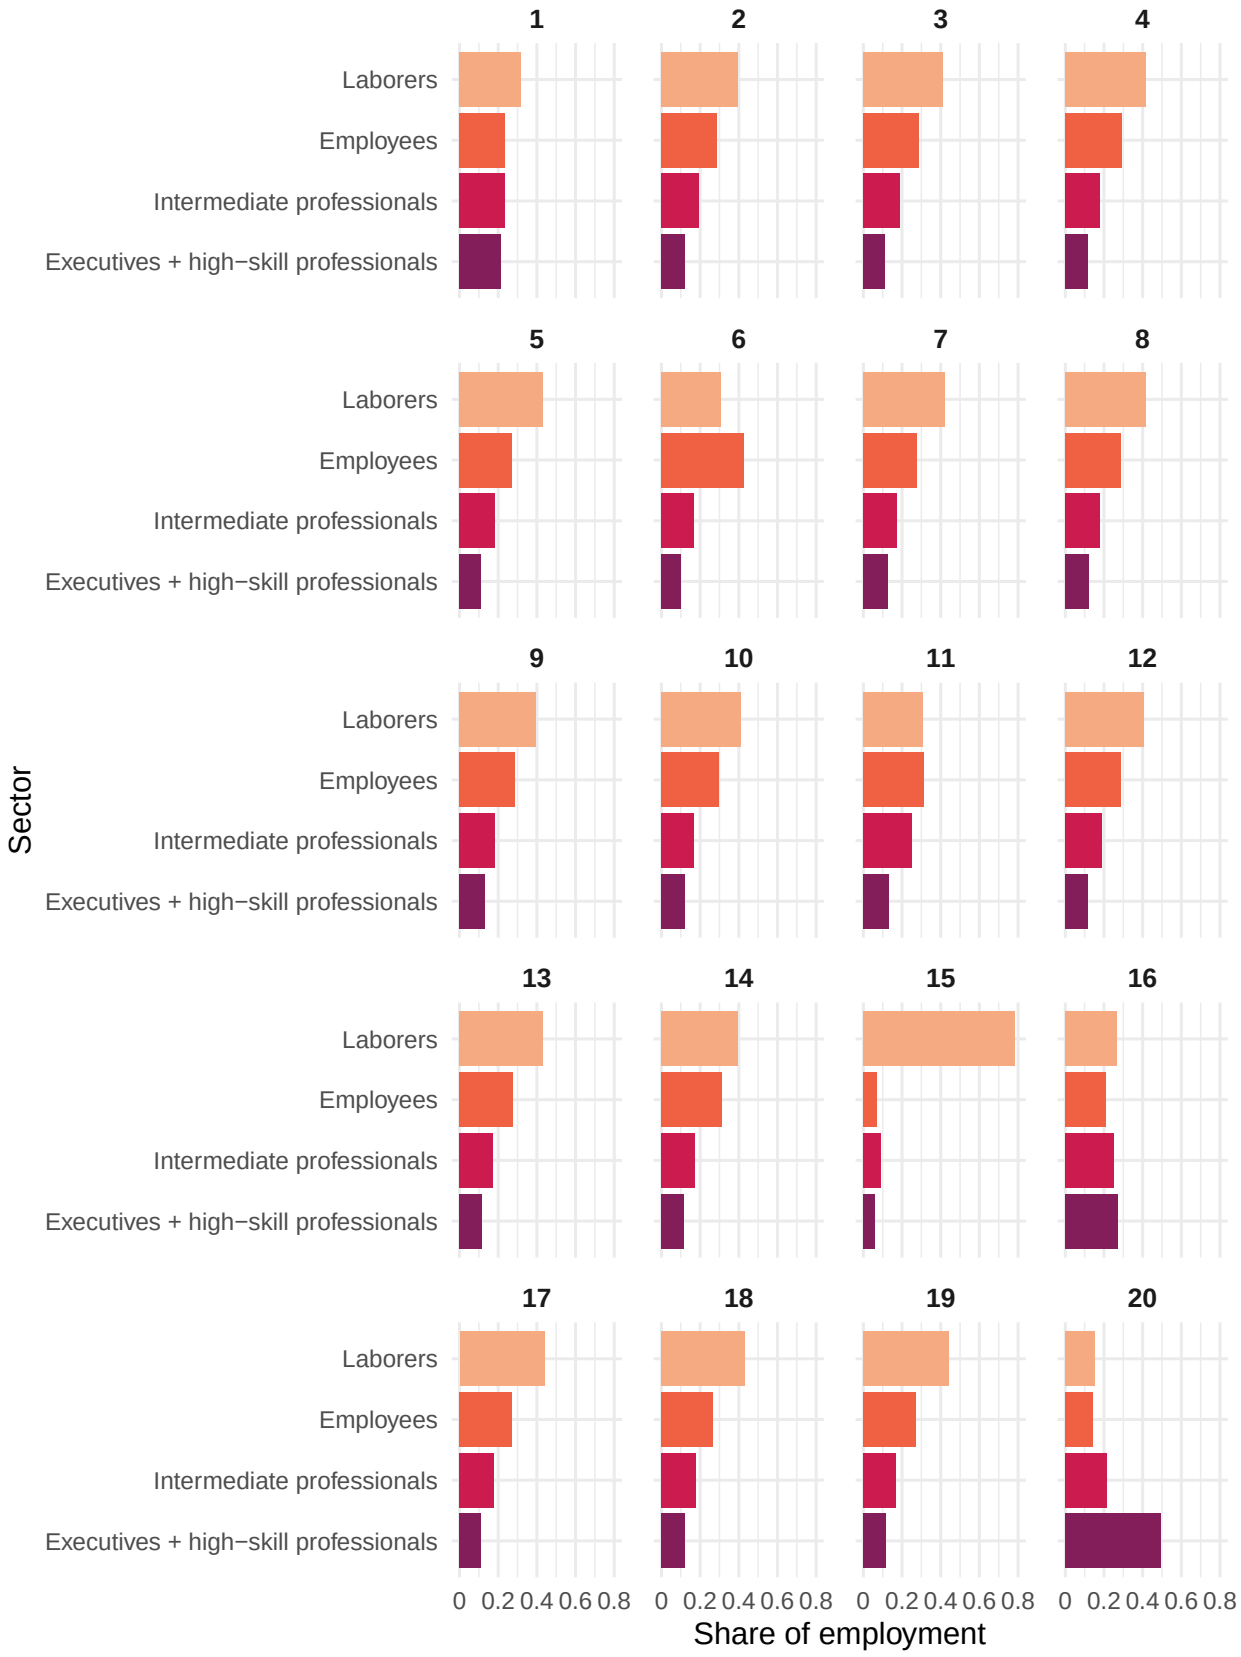


FIGURE D.4. Geographic Composition of Data-Driven Labor Markets

