The Two-Body Problem: Dual-Earner Job Search, Migration, and Labor Market Outcomes

Hanno Foerster* Robert Ulbricht†

May 21, 2026

Abstract

We develop a spatial directed search model of dual-earner households in which spouses choose where to search. Because households are the unit of mobility but workers are the unit of matching, the model naturally features two colocation constraints: *geographic misalignment* in spouses' earnings prospects and a *coordination friction*, because moving without job loss requires a double coincidence of offers. Estimated on U.S. data, these constraints amplify gender gaps in employment and wages, reduce migration, especially among power couples, and create persistent trailing-spouse losses borne mainly by women. Remote work mitigates these losses by making jobs portable across locations.

Keywords: Dual-earner migration, gender inequality, household job search, labor misallocation, remote work, spatial directed search.

JEL Classification: E24, J16, J61, J64, R23.

*Boston College, hanno.foerster@bc.edu
†Boston College, ulbricht@bc.edu

1 Introduction

A substantial share of U.S. workers belong to dual-earner households.[1] These households face a "two-body problem": they must secure employment for two workers in a location that serves both careers. This creates two distinct constraints. First, locations that offer strong earnings prospects for one spouse may not do so for the other. Second, even when a location is attractive for both spouses, relocating without job loss requires two offers in the same place. We call the first constraint *geographic misalignment* and the second the *coordination friction*.

This paper develops a spatial search model of dual-earner households to formalize and quantify these constraints. Building on Menzio & Shi (2010, 2011), search is directed across locations, so households actively choose where to search and migrate. This makes geographic misalignment consequential for search behavior. At the same time, hiring is decentralized. Employers make offers to individual workers, and offers arrive independently across spouses. This lack of employer coordination gives rise to the coordination friction.

We use the framework to study three related questions. First, we ask how pre-existing gender differences in productivity, search efficiency and job separation rates translate into gender gaps in employment, earnings and post-migration outcomes. Second, we quantify how geographic misalignment and the coordination friction shape migration rates, long-run sorting across locations, and earnings. Third, we study how remote work changes these outcomes by making some jobs portable across locations, thereby allowing households to relocate without requiring both spouses to obtain new jobs in the same place.

Quantification We calibrate our framework to the U.S. labor market. Households differ by their members' employment status, occupations, human capital, and children. Commuting zones differ in amenities, rents, childcare costs, and gender- and occupation-specific productivities. We calibrate these differences, combining household data from the American Community Survey (ACS) and the Current Population Survey (CPS) with geolocation data, commuting zone data from the Opportunity Atlas, and web-scraped data to inform the geography of the U.S. labor market. We further match occupation-level job and migration flows, the steepness of job-ladders, and the cross-sectional distribution of households over commuting zones.

Results The quantitative model reproduces observed gender gaps in wages and employment without explicitly targeting them during the estimation. Specifically, men are 27% more likely to be employed and have 31% higher wages than women. Following migration these gaps widen to 45% and 41%. Within the model, the gaps and their widening can be decomposed

[1] We define dual-earner households as married or cohabiting couples with both spouses in the labor force. In the 2010–2019 American Community Survey, they represent 66% of all such couples.

into four gender-specific inputs: differences in productivity levels, productivity dispersion, job separation rates, and search efficiency. We provide an exact decomposition, finding that differences in productivity levels account for the largest share, followed by job separation rates.[2]

The colocation constraints amplify pre-existing gender gaps along three margins: they shape households' destination choices, affect the timing of relocation, and alter spouses' labor-market trajectories after a move. Geographic misalignment distorts household sorting across locations, compromising women's earnings prospects more than men's. The coordination friction creates trailing spouses, with persistent losses in human capital and wages. These losses fall disproportionately on women, who are the trailing spouse in four out of five work-related moves.

Our model clarifies when these constraints bind. Geographic misalignment matters when household location choices are driven by earnings prospects and when there are no gains from diversifying search across locations. The coordination friction binds when the household value of receiving two colocated offers exceeds the sum of receiving either offer separately—that is, when the value function is locally convex in the number of colocated offers. This condition is more likely to hold when both spouses are employed and childless, job ladders are steep, migration costs are large, and when the gains from diversifying search across locations are large.

Quantitatively, the coordination friction accounts for the larger losses. Geographic misalignment lowers average earnings by 1.2% for women and 0.5% for men, mostly through its effect on household sorting across locations. By contrast, eliminating the coordination friction raises migration by 13% in the long run, increases women's lifetime earnings gains from migration by 19%, and raises average earnings by 2.7% for women and 0.9% for men. Its effects are largest among "power couples" with two employed spouses and above-average human capital.

Finally, we examine remote work as a relaxation of the two-body problem. Remote jobs are portable across locations, allowing a spouse to remain employed when the household moves and thereby mitigating the coordination friction. We simulate the introduction of remote work by allowing jobs to convert to remote positions at occupation-specific rates calibrated to match the distribution of remote-workable jobs. Migration initially declines as households wait for remote conversion, then gradually recovers as remote work spreads. Employment and earnings rise, especially for women. In the long run, remote work raises earnings by up

[2]In our framework, gender differences in productivity may capture a variety of underlying mechanisms that lead women's wages to diverge from men's within occupations, including gender discrimination, differences in bargaining power, compensating differentials associated with workplace amenities, and differences in working hours (see, e.g., Blau & Kahn 2017).

to 1.4% for women and 0.4% for men, offsetting a substantial share of the earnings losses generated by constraints on dual-earner mobility.

Related literature The paper's first contribution is to macro-labor search theory. We introduce dual-earner households into a spatial directed-search environment. The framework builds on directed search models by Menzio & Shi (2010, 2011), Menzio et al. (2016), Schaal (2017), and Herkenhoff et al. (2023). The key departure is that the household is the unit of mobility, while the worker remains the unit of matching. This creates a household-level matching externality: a nonlocal offer that is privately attractive for a worker can reduce the earnings prospects of their spouse and destroy the spouse's current match, reflecting precisely the two colocation constraints at the core of this paper. On the computational side, we benefit from block recursivity (Menzio & Shi 2010; Wright et al. 2021), which makes our quantitative model tractable despite a large state and action space.

Second, we contribute to the literature on dual-earner migration decisions, dating back to Mincer (1978).[3] This literature has explored how migration decisions vary with household composition and gender-specific labor-market opportunities (Costa & Kahn 2000; Lundberg & Pollak 2003; Guler et al. 2012; Lessem 2017; Braun et al. 2021; Alonzo et al. 2023; Gemici 2023; Glaeser et al. 2025; Rueda & Wilemme 2025), labor-market policy (Venator 2025), and the geographic distribution of occupations (Román 2025). Relative to this work, we embed dual-earner migration in a directed spatial search model in which households choose where each spouse searches, while job offers arrive separately to individual workers. This combination lets us isolate the two constraints at the core of the paper.[4] Quantitatively, we estimate the model at the commuting-zone level, capturing realistic heterogeneity in rents, amenities, and earnings prospects that is typically averaged out in coarser spatial frameworks and is central to measuring geographic misalignment.[5]

Third, the paper speaks to work on labor misallocation. Prior work has examined misallocation arising from search frictions (Hosios 1990), monopsony power (Galenianos et al. 2011; Rabinovich & Wolthoff 2022), spatial frictions (Şahin et al. 2014; Findeisen et al. 2021; Kuhn et al. 2021; Zerecero 2021; Azkarate-Askasua et al. 2023; Jung et al. 2023),

[3]This literature is at the intersection of a somewhat larger literature studying dual-earner job search (e.g., Dey & Flinn 2008; Pilossoph & Wee 2021; Flabbi & Mabli 2018) and the literature on domestic migration, which typically focuses on single-earners (e.g., Coen-Pirani 2010; Kennan & Walker 2011; Kaplan & Schulhofer-Wohl 2017; Coen-Pirani & Sieg 2019; Piyapromdee 2020).

[4]At a more technical level, directed search prevents the double-coincidence problem from mechanically scaling with the number of locations. Spouses can target the same destinations, so the coordination friction reflects independently arriving offers rather than random offer dispersion across space.

[5]Existing models of dual-earner household migration are typically calibrated to a small number of locations, such as two locations, nine Census divisions, or 48 mainland U.S. states. Meaningfully capturing geographic misalignment requires finer spatial heterogeneity, motivating our quantification at the commuting-zone level.

and information frictions (Jovanovic 1979, 1984; Baley et al. 2022). Our paper introduces constraints on dual-earner mobility as a source of labor misallocation.

Finally, we contribute to the emerging literature on remote work and its labor market implications (Alon et al. 2020; Bagga et al. 2025; Berresheim 2025; Faberman et al. 2026), by quantifying how remote work mitigates constraints on dual-earner mobility and reshapes migration and earnings dynamics for dual-earner households.

Roadmap The paper proceeds as follows. Section 2 presents the model, and Section 3 formalizes the two colocation constraints. Section 4 describes the calibration. Section 5 decomposes gender gaps, Section 6 quantifies the effects of the colocation constraints, and Section 7 studies remote work as a relaxation of these constraints. Section 8 concludes. All proofs are in Appendix A.

2 General Framework

We develop a spatial search model of the labor market for dual-earner households that formalizes the two-body problem. The model embeds two colocation constraints. First, directed search gives households an active choice over where to search and migrate, making *geographic misalignment* in spouses' location-specific earnings prospects consequential for search behavior. Second, because hiring is decentralized and offers arrive independently across spouses, households face a *coordination friction*: if one spouse accepts a nonlocal job, the household moves and the other spouse's job is destroyed.

2.1 Environment

Preferences and technology Time is continuous and extends forever. There is a finite set of locations, indexed by $r \in \mathcal{R}$. The economy is populated by an endogenous measure of one-vacancy firms and a unit measure of households. Each household consists of two adult workers (or "spouses"), indexed by $i \in \{1, 2\}$, and zero or more children. While we will equate i with workers' "gender" in our quantification, there is no need at this stage to make assumptions about the gender-composition within households. Firms and households are risk neutral and share the same effective discount rate ρ.

Following the literature, we assume that search and matching is privately efficient. In order to characterize labor and migration flows, it then suffices to specify the sum of a household's

instantaneous utility flow and the labor product of its employed members.[6] Let

$$u(\boldsymbol{e}, \boldsymbol{s}, r) = \bar{u}(\boldsymbol{e}, \boldsymbol{s}, r) + \sum_{i \in \{1,2\}} z_i(\boldsymbol{s}, r) \cdot \mathbb{1}_{e_i=1} \tag{1}$$

denote this joint value flow. Here, $\boldsymbol{e} \equiv (e_1, e_2) \in \{0,1\}^2$ is the employment status of the household's adult members, $\bar{u}$ is their utility flow net of earnings, $z_i(\boldsymbol{s}, r)$ is the location-specific labor product of spouse i, $r \in \mathcal{R}$ is the household's current location, and $\boldsymbol{s} \in \mathcal{S}_1 \times \cdots \times \mathcal{S}_{n_s}$ is a generic "catch-all" state that captures other transitory and persistent characteristics of the household. For example, in our quantification, $\boldsymbol{s}$ includes the occupation of both spouses, their human capital, and whether or not they have children. We assume that $\boldsymbol{s}$ has finite support $\mathcal{S}_k$ in all its dimensions $k \in \{1, \ldots, n_s\}$.

Other than through search and migration, a household's type $(\boldsymbol{e}, \boldsymbol{s}, r)$ evolves stochastically with Poisson arrival rates given by $\pi(\boldsymbol{e}', \boldsymbol{s}', r' | \boldsymbol{e}, \boldsymbol{s}, r)$. In our quantification, we specify π to expose households to exogenous job separations, human capital dynamics, location preference shocks, and the arrival and departure of children.

Labor markets and migration The labor market is organized in a continuum of submarkets indexed by the location of jobs $q \in \mathcal{R}$, the worker type $(i, \boldsymbol{s})$, and the firm's share y of the joint value of the match.[7] Workers direct their search toward these submarkets, choosing both a firm share y and a search effort $\kappa_{i,q}$ for each location q. Specifically, each spouse $i \in \{1,2\}$ is endowed with a type-specific search budget $\bar{\kappa}_i(\boldsymbol{e}, \boldsymbol{s})$, which they can allocate to search across submarkets in different locations subject to

$$\left(\sum_{q \in \mathcal{R}} \kappa_{i,q}^{\frac{1+\eta}{\eta}} \right)^{\frac{\eta}{1+\eta}} \leq \bar{\kappa}_i(\boldsymbol{e}, \boldsymbol{s}). \tag{2}$$

Here, $\eta \geq 0$ is the elasticity of substitution between locations, governing the spatial "breadth" of search. In the limit where $\eta \to \infty$, workers allocate their entire search budget to the single location with the highest gains from search, as is usually the case in directed search models. At the other extreme where $\eta = 0$, diversifying search is costless and workers allocate $\bar{\kappa}_i(\boldsymbol{e}, \boldsymbol{s})$ units of search effort to each location with positive search gains as in the literature on multiple job applications (Albrecht et al. 2006; Kircher 2009; Galenianos & Kircher 2009).

[6] Given private efficiency and risk-neutrality, wages only split the surplus and do not affect search and migration choices. Hence, we can work directly with the *joint* value flow $u(\cdot)$, deferring value-splitting to below.

[7] Indexing submarkets by firms' share y is equivalent to indexing by a utility level for the household along with households' employment status $\boldsymbol{e}$. While indexing by y allows us to eliminate $\boldsymbol{e}$ from the submarket partition, the elimination is without loss as long as vacancies are created at constant returns to scale.

Vacancies are created by an infinite supply of potential firms, which can open vacancies in any submarket at flow costs c. Vacancies and workers in a given submarket come together through a frictional matching process. In particular, a worker that allocates search effort $\kappa_{i,q}$ to submarket $\psi \equiv (q, y, i, \boldsymbol{s})$ meets a vacancy at rate $\kappa_{i,q}\lambda(\theta_\psi)$, where θ_ψ is the ratio between vacancies posted and the measure of workers' combined efforts in submarket ψ. Similarly, a vacancy posted in submarket ψ meets a worker at rate $\lambda(\theta_\psi)/\theta_\psi$. As usual, we assume that the contact function λ is twice differentiable, strictly increasing and concave, with $\lambda(0) = \lambda'(\infty) = 0$ and $\lambda'(0) = \infty$.

When a worker meets a firm, the firm offers a wage contract with present discounted value equal to the match value net of the firm's share y and hires the worker. Following Menzio & Shi (2010, 2011), we assume that the underlying contract space is complete. In particular, endogenous separations and on-the-job search maximize the joint value of the household and all its current employers.

A hire entails migration whenever the new job is in a location q that differs from a household's current location r. In this case, due to the coordination friction, the spouse without the job offer quits their job and the household moves to q. Migration entails a utility cost $\chi(q|\boldsymbol{s}, r)$, normalized so that $\chi(r|\cdot, r) = 0$.

Discussion Three remarks are in order. First, while our analysis centers on dual-earner households, the model naturally generates labor force exits — non-employed individuals who exert zero search effort. In our calibration, women in households with children have the highest non-participation rate, at 18%, while exits among men are rare.[8]

Second, our model assumes that households reside in a single location. This assumption is motivated by the low prevalence of long-distance relationships in the United States.[9] A potential concern is that such relationships might become more common around migration events. To examine this, we use data from the Panel Study of Income Dynamics and find no evidence that the frequency of long-distance relationships increases either before or after migration (see Appendix Figure A.I).

Third, one might ask whether households can secure job offers for both spouses simultaneously, either due to employer coordination or due to the possibility of storing offers over short horizons. Our baseline model rules out these possibilities, reflecting precisely the presence of the coordination friction. Section 4.3 shows that the baseline model matches post-migration dual-employment rates. Complementing evidence from DellaVigna et al. (2021) and Cortés et al. (2023), this inference suggests that such forms of employer coordination and offer storage

[8]The non-participation rate among women is close to that of prime-age partnered women in our 2010-2019 ACS sample, which is 23%.

[9]The share of married couples living apart between 2010 and 2019 is 0.026 in the CPS.

are limited in the data.

2.2 Equilibrium Characterization

Vacancy creation By free entry, the value of creating a vacancy must be zero in every submarket. That is,

$$c\theta_\psi = \lambda(\theta_\psi)\max\{0, y\}, \tag{3}$$

which pins down the market tightness θ_ψ as a function of y.

Search, migration, and separation Next, consider the search, migration and separation policies of workers. Private efficiency implies that the policies maximize the joint value of a household and its current employers, given by

$$\begin{aligned}\rho V(\boldsymbol{e}, \boldsymbol{s}, r) = u(\boldsymbol{e}, \boldsymbol{s}, r) &+ \sum_{\boldsymbol{e}', \boldsymbol{s}', r'} \pi(\boldsymbol{e}', \boldsymbol{s}', r' | \boldsymbol{e}, \boldsymbol{s}, r)\Big(V(\boldsymbol{e}', \boldsymbol{s}', r') - V(\boldsymbol{e}, \boldsymbol{s}, r)\Big) \\ &+ \max_{\{\kappa_{i,q}, y_{i,q}\}} \sum_{i,q} \kappa_{i,q}\lambda(\theta_{i,q})\Big(V(\boldsymbol{e}^{\text{post},i}, \boldsymbol{s}, q) - y_{i,q} - \chi(q|\boldsymbol{s}, r) - V(\boldsymbol{e}, \boldsymbol{s}, r)\Big) \\ &+ \lim_{\epsilon \to \infty} \epsilon \sum_i \max\Big\{0, V(\boldsymbol{e}^{\text{sep},i}, \boldsymbol{s}, r) - V(\boldsymbol{e}, \boldsymbol{s}, r)\Big\}. \end{aligned} \tag{4}$$

The joint value is composed of four terms: (i) the current value flow defined in (1); (ii) value changes induced by exogenous shocks to $(\boldsymbol{e}, \boldsymbol{s}, r)$; (iii) the value of search by either spouse, maximized subject to the θ–y frontier posed by (3); and (iv) the option for either spouse to endogenously quit their job, which they do at rate $\epsilon \to \infty$ whenever the surplus of their current job is negative. Here, $\boldsymbol{e}^{\text{post},i}$ is a household's employment status after spouse i accepts a new job, defined by $e_i^{\text{post},i} = 1$ and $e_{-i}^{\text{post},i} = e_{-i} \cdot \mathbb{1}_{q=r}$ where "$-i$" denotes the spouse without a job offer. Similarly, $\boldsymbol{e}^{\text{sep},i}$ is the employment status after spouse i quits their job, defined by $e_i^{\text{sep},i} = 0$ and $e_{-i}^{\text{sep},i} = e_{-i}$.

It remains to characterize the optimal choice of $\{\kappa_{i,q}, y_{i,q}\}$ in households' job search. Consider first the value split between households and firms. From (3), the market tightness is increasing in firms' share y, creating a trade-off for workers to search in submarkets with higher job-finding rates versus searching in submarkets where they retain a larger share of the surplus. Maximizing the third term in (4) subject to (3), the optimal market tightness is given by

$$\theta_{i,q}(\boldsymbol{e}, \boldsymbol{s}, r) = \lambda'^{-1}\left(\frac{c}{V(\boldsymbol{e}^{\text{post},i}, \boldsymbol{s}, q) - \chi(q|\boldsymbol{s}, r) - V(\boldsymbol{e}, \boldsymbol{s}, r)}\right), \tag{5}$$

which in turn pins down the matching rate per unit of search effort, $\lambda(\theta_\psi)$.

Finally, consider the allocation of search effort across locations, $\{\kappa_{i,q}\}$. Let $\Lambda_{i,q}$ denote

spouse i's expected gain per unit of search effort allocated to location q,

$$\Lambda_{i,q}(\boldsymbol{e},\boldsymbol{s},r) \equiv \lambda(\theta^*_{i,q})\Big(V(\boldsymbol{e}^{\text{post},i},\boldsymbol{s},q) - \chi(q|\boldsymbol{s},r) - V(\boldsymbol{e},\boldsymbol{s},r)\Big) - \theta^*_{i,q}c,$$

with $\theta^*_{i,q} \equiv \theta_{i,q}(\boldsymbol{e},\boldsymbol{s},r)$ denoting the optimal market tightness in (5). Maximizing the joint value (4) subject to the search budget (2), then yields

$$\kappa_{i,q}(\boldsymbol{e},\boldsymbol{s},r) = \Lambda_{i,q}(\boldsymbol{e},\boldsymbol{s},r)^{\eta}\left(\sum_{v}\Lambda_{i,v}(\boldsymbol{e},\boldsymbol{s},r)^{1+\eta}\right)^{-\frac{\eta}{1+\eta}}\bar{\kappa}_i(\boldsymbol{e},\boldsymbol{s}). \tag{6}$$

Together with (5) this pins down job-finding and migration rates,

$$f_{i,q}(\boldsymbol{e},\boldsymbol{s},r) \equiv \kappa_{i,q}(\boldsymbol{e},\boldsymbol{s},r)\cdot\lambda(\theta_{i,q}(\boldsymbol{e},\boldsymbol{s},r)), \tag{7}$$

completing the characterization of search policies.

Implications for job-finding rates Equations (5) to (7) imply that job-finding rates are endogenous: households choose both where to search and which submarket tightness to target through the value split. As search gains vary with the household's state and location, so do job-finding rates. In particular, when migration costs make local search more attractive than nonlocal search, a trailing spouse who moves to a new location searches more intensively there and targets a tighter submarket, thereby raising the local job-finding rate.

Steady state equilibrium In this economy, all policy rules are functions of only the idiosyncratic household type $(\boldsymbol{e},\boldsymbol{s},r)$. An equilibrium is a collection of maps from $(\boldsymbol{e},\boldsymbol{s},r)$ to search and separation policies satisfying (4) to (6) along with a value split satisfying (3). Unless otherwise mentioned, we focus on the case where the cross-sectional distribution over $(\boldsymbol{e},\boldsymbol{s},r)$ is at steady state.[10]

3 Colocation Constraints

The framework provides a general setting for analyzing job search and migration in dual-earner households. In this section, we formalize two colocation constraints: *geographic misalignment* in spouses' earnings prospects across locations, and a *coordination friction* arising from independently arriving job offers. The first is a spatial trade-off; the second is a matching friction. This section characterizes when each constraint matters. Section 6 quantifies how

[10]The cross-sectional distribution is characterized by a standard Kolmogorov forward equation, detailed in Appendix A.1.

they affect labor market outcomes, both conditional on relocation and through their effects on mobility and long-run sorting across locations.

3.1 Geographic misalignment

Locations can rank differently for each spouse's earnings prospects, creating trade-offs in dual-earner migration decisions. We refer to these rank differences as *geographic misalignment*. When migration is primarily driven by earnings considerations, misalignment can reduce the attainable gains from relocating. We formalize this idea with the following benchmark.

Geographic alignment benchmark Rank commuting zones by the productivities $z_i(\boldsymbol{s}, r)$ for each state $\boldsymbol{s}$ and each spouse i, and let $rank_{i,\boldsymbol{s}}(r)$ denote the productivity rank of location r.[11] Consider a counterfactual reshuffling of spouse 1's productivities that aligns the rank of each commuting zone with that of spouse 2:

$$z_1^{aligned}(\boldsymbol{s}, r) = z_1\left(\boldsymbol{s}, rank_{1,\boldsymbol{s}}^{-1}\left(rank_{2,\boldsymbol{s}}(r)\right)\right).$$

The benchmark aligns spouses' rankings over earnings prospects while keeping the marginal earnings distribution across locations fixed. It thereby isolates the cost of having one household location determine the career opportunities for both spouses.

To begin, note that aligning spouses' earnings prospects clearly maximizes the earnings frontier available to dual-earner households.

Lemma 1. *Fix $\boldsymbol{s}$ and let*

$$\mathcal{Z}(\boldsymbol{s}) = \max_{q \in \mathcal{R}} \{z_1(\boldsymbol{s}, q) + z_2(\boldsymbol{s}, q)\}$$

denote the household's earnings frontier across locations. Among all permutations of $\{z_1(\boldsymbol{s}, r)\}_{r\in\mathcal{R}}$ holding $\{z_2(\boldsymbol{s}, r)\}_{r\in\mathcal{R}}$ fixed, $\mathcal{Z}(\boldsymbol{s})$ is maximized by the geographic alignment benchmark.

If earnings potentials were the only source of spatial heterogeneity, then maximizing the earnings frontier would also map into higher realized post-migration earnings and utilities. The next proposition formalizes this.

Proposition 1. *Suppose $\eta \to \infty$ and $\bar{u}(\cdot, \cdot, r)$, $\pi(\cdot, \cdot, \cdot|\cdot, \cdot, r)$ and $\chi(q|\cdot, r)$ are independent of q and r. Fix $\boldsymbol{e}, \boldsymbol{s}, r$ and a spouse i such that the state entails migration led by i. Then the*

[11] Assume ranks are unique; ties can be broken arbitrarily.

post-migration continuation value is

$$V^{\text{post},i}(\boldsymbol{e}, \boldsymbol{s}) = \max_{q \in \mathcal{R}} V\left(\boldsymbol{e}^{\text{post},i}, \boldsymbol{s}, q\right). \tag{8}$$

Moreover, among all permutations of $\{z_1(\boldsymbol{s}, r)\}_{r \in \mathcal{R}}$ *holding* $\{z_2(\boldsymbol{s}, r)\}_{r \in \mathcal{R}}$ *fixed, post-migration values are maximized pointwise by the geographic alignment benchmark.*

Outside the conditions of Proposition 1, geographic alignment need not be beneficial. If η is finite, households optimally diversify search across multiple destinations. In that case, aligning productivity ranks can lower the attractiveness of secondary destinations.[12] By the same logic, when migration is driven by non-wage considerations—such as amenities, cost-of-living, differential exposure to shocks, or heterogeneity in migration costs—the optimal destination need not maximize earnings potential, so aligning productivity ranks need not raise (and may even lower) realized migration gains.

In our quantitative model, migration decisions and realized gains reflect these additional sources of spatial heterogeneity. Section 6.1 studies how geographic misalignment affects equilibrium outcomes once households' search, migration, and long-run sorting respond endogenously to these additional forces.

3.2 Coordination friction

Dual-earner households seeking to relocate without unemployment require a double coincidence of job offers. In our baseline economy, offers arrive independently across spouses, reflecting the absence of employer coordination to resolve this coincidence. We refer to this lack of coordination as *coordination friction*.

The following benchmark analyzes when the friction binds. To do so, it gives households the option to correlate their job offers, eliminating the double-coincidence problem. We then infer the friction to bind in any state where households would prefer offers to be correlated rather than independent.

Correlated Matching Benchmark The correlated matching benchmark replicates the baseline economy but allows spouses to receive joint job offers at a destination-specific rate ω_q. Consider two spouses that exert search efforts $\{\kappa_{i,q}\}_{i=1,2}$ to find jobs in submarkets $\{\psi_{i,q}\}_{i=1,2}$ in some location q. Given ω_q, the spouses obtain a joint job offer in $\psi_{1,q}$ and $\psi_{2,q}$ at rate ω_q, and obtain individual offers at the residual rates $\kappa_{i,q}\lambda(\theta_{\psi_{i,q}}) - \omega_q$.

[12] Specifically, migration becomes less attractive to destinations q with $rank_{1,\boldsymbol{s}}(q) > rank_{2,\boldsymbol{s}}(r)$.

The benchmark is characterized by a joint assignment of correlation rates, search efforts, and submarkets to households,

$$\mathcal{C} : (\boldsymbol{e}, \boldsymbol{s}, r, V) \mapsto \{\omega_q\}_{q\in\mathcal{R}}, \{\kappa_{i,q}, \psi_{i,q}\}_{i,q\in\{1,2\}\times\mathcal{R}}.$$

Crucially, households are indexed by a continuation value function V in addition to their current state, enabling us to apply the benchmark to the steady state of the baseline economy as well as any transition path. The benchmark assignment $\mathcal{C}$ maximizes households' lifetime utility subject to the feasibility constraint

$$0 \leq \omega_q \leq \min_{i\in\{1,2\}} \left\{\kappa_{i,q}\lambda(\theta_{\psi_{i,q}})\right\} \quad \text{for all } q \in \mathcal{R}. \tag{9}$$

We note that the benchmark does not change the primitives of the search technology. That is, workers face the same constraint (2) on their allocation of search effort $\{\kappa_{i,q}\}$, firms face the same cost of vacancy creation c, and matching is subject to the same frictional matching function λ. As a result, both economies share the same feasibility frontier for *total matches*, differing only in how those matches can be distributed across households.

We say that for a household of type $(\boldsymbol{e}, \boldsymbol{s}, r)$ with continuation value V, the coordination friction binds for job searches in location q if and only if matching would optimally be correlated in the counterfactual economy, $\omega_q(\boldsymbol{e}, \boldsymbol{s}, r, V) > 0$. The following proposition characterizes when this is the case.

Proposition 2. *Fix a continuation value function V. Let*

$$\Delta V_{i,q}(\boldsymbol{e}, \boldsymbol{s}, r) = V(\boldsymbol{e}^{\text{post},i}, \boldsymbol{s}, q) - \chi(q|\boldsymbol{s}, r) - V(\boldsymbol{e}, \boldsymbol{s}, r)$$

denote the gains from an individual match by spouse i. Similarly, let

$$\Delta V_q^{\text{corr}}(\boldsymbol{e}, \boldsymbol{s}, r) = V(\boldsymbol{e}^{\text{corr}}, \boldsymbol{s}, q) - \chi(q|\boldsymbol{s}, r) - V(\boldsymbol{e}, \boldsymbol{s}, r)$$

denote the gains from a correlated match, where $\boldsymbol{e}^{\text{corr}} = (1, 1)$. Then correlated matching is beneficial if and only if

$$\Delta V_q^{\text{corr}}(\boldsymbol{e}, \boldsymbol{s}, r) > \sum_{i\in\{1,2\}} \Delta V_{i,q}(\boldsymbol{e}, \boldsymbol{s}, r). \tag{10}$$

If condition (10) holds, the correlated matching rate is optimally set to its upper bound, $\omega_q = \min_{i\in\{1,2\}} \left\{\kappa_{i,q}\lambda(\theta_{\psi_{i,q}})\right\}$. Otherwise, it is optimally set to its lower bound, $\omega_q = 0$.

Intuitively, as illustrated in Figure 1, condition (10) evaluates the curvature of $V - \chi$ with respect to the number of colocated job offers. If this mapping is convex, the household

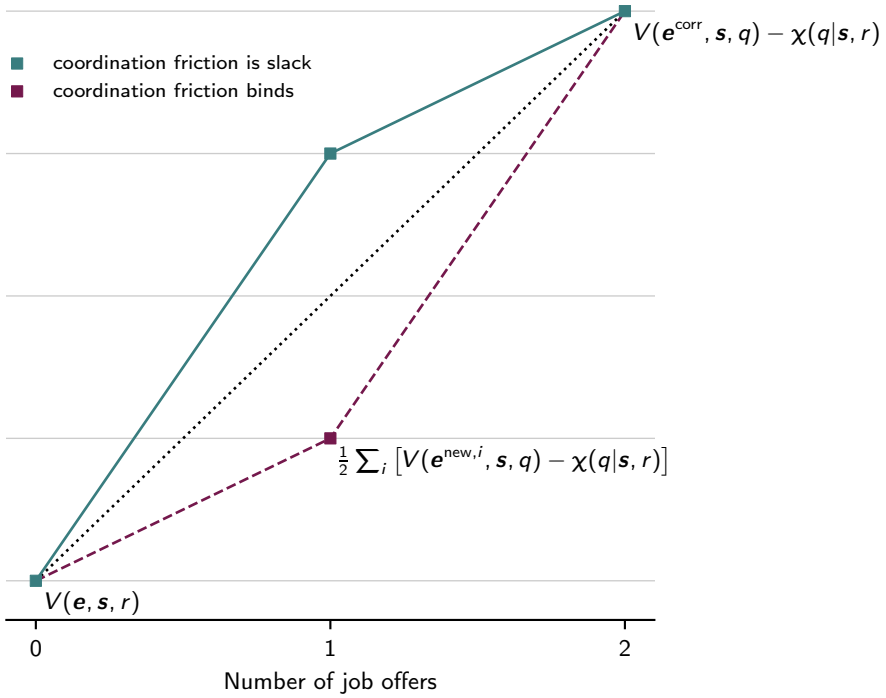

Figure 1: Illustration of condition (10). The coordination friction binds if and only if, averaged across spouses, $V - \chi$ is convex in the number of colocated job offers as in the dashed graph, and is slack if $V - \chi$ is concave as in the solid graph.

benefits from increased risk through correlated job offers, implying that the coordination friction is binding. Conversely, if the mapping is concave, the household is better off hedging by keeping job offers statistically independent, rendering the coordination friction slack.

Simple Example Before quantifying our framework, we briefly explore a simple example to illustrate the economic forces that determine the curvature of V. Consider a household with initial type $(\boldsymbol{e}_0, \boldsymbol{s}_0, r_0)$. There are no exogenous shocks, $\pi(\cdot|\cdot) = 0$, and both spouses and all locations $q \neq r_0$ are symmetric. Specifically, we have $\bar{\kappa}_i(\boldsymbol{e}, \boldsymbol{s}) = \bar{\kappa}$, $\chi(q|\boldsymbol{s}, r) = \chi$, and

$$u(\boldsymbol{e}, \boldsymbol{s}, r) = \begin{cases} u_0 & r = r_0 \\ u_{e_1+e_2} & r \neq r_0, \end{cases}$$

with $u_0 + \rho\chi \leq u_1 \leq u_2$. Finally, we simplify by considering the limit where the matching function is inelastic in vacancies, $\lambda(\theta) = \lim_{\gamma \to 0} \theta^\gamma$.

The example admits a simple recursive solution where $V(\boldsymbol{e}, \boldsymbol{s}, r) \in \{V_0, V_1, V_2\}$. Specifically, if $r \neq r_0$ and both spouses are employed, we have $\rho V_2 = u_2$. If $r \neq r_0$ and only one spouse is

employed, we have $\rho V_1 = u_1 + \bar{\kappa}\,(V_2 - V_1)$. Finally, the initial value at $r = r_0$ is given by

$$\rho V_0 = u_0 + 2\bar{\kappa} N^{\frac{1}{1+\eta}}\,(V_1 - \chi - V_0)\,,$$

where N is the number of locations $q \neq r_0$.

Intuition Evaluating condition (10) for the simple example, the coordination friction binds if and only if

$$\frac{u_0 + \Omega(N^{\frac{1}{1+\eta}})}{2} + \frac{u_2 - \rho\chi}{2} \geq u_1 - \rho\chi \tag{11}$$

with

$$\Omega(x) = 2\bar{\kappa}\left(\frac{x-1}{\rho + 2\bar{\kappa}x}\right)\left(u_1 - u_0 - \rho\chi + \frac{\bar{\kappa}}{\rho+\bar{\kappa}}\,(u_2 - u_1)\right).$$

The condition reveals several distinct factors affecting the curvature of $V - \chi$, illustrated in Figure 2. First, consider the case where there are no migration costs, $\chi = 0$, and search effort is infinitely elastic across locations, $\eta \to \infty$. In this case, condition (11) simplifies to $(u_0 + u_2)/2 \geq u_1$, so the curvature of $V - \chi$ in the number of simultaneous job offers is entirely determined by the curvature of u.

In particular, given u_1 and u_2, V is more convex for larger u_0, reflecting that the coordination friction is more likely to bind for households with a high initial value flow such as when both spouses are employed (illustrated in Panel a of Figure 2). Conversely, the coordination friction is more likely to be slack for households with a low initial value flow such as when both spouses are non-employed (Panel b).

Similarly, given u_0 and u_2, V is more convex the smaller u_1 (Panel d). Intuitively, this captures, in reduced form, economies where falling off the job ladder comes with high costs; e.g., due to human capital depreciation (e.g., Pissarides 1992; Ljungqvist & Sargent 1998; Jung & Kuhn 2019), deteriorated bargaining positions (e.g., Postel–Vinay & Robin 2002; Cahuc et al. 2006; Lise & Robin 2017; Bagger & Lentz 2019), or slippery bottom job rungs (Jarosch 2023). All of these mechanisms increase the cost of becoming a trailing spouse, making it more likely that the coordination friction binds. Conversely, factors reducing the value flow of dual employment u_2, such as having children when childcare costs are large, reduce the convexity of V and make it less likely that the coordination friction binds (Panel c).

Next, consider the case where migration costs are strictly positive, $\chi > 0$. As illustrated in Panel (e), this trivially introduces convexity in the search gains, making it more likely that the coordination friction binds. Intuitively, this is because migration costs accrue regardless of whether the household has one or two job offers at hand, which raises the relative returns of having two job offers.

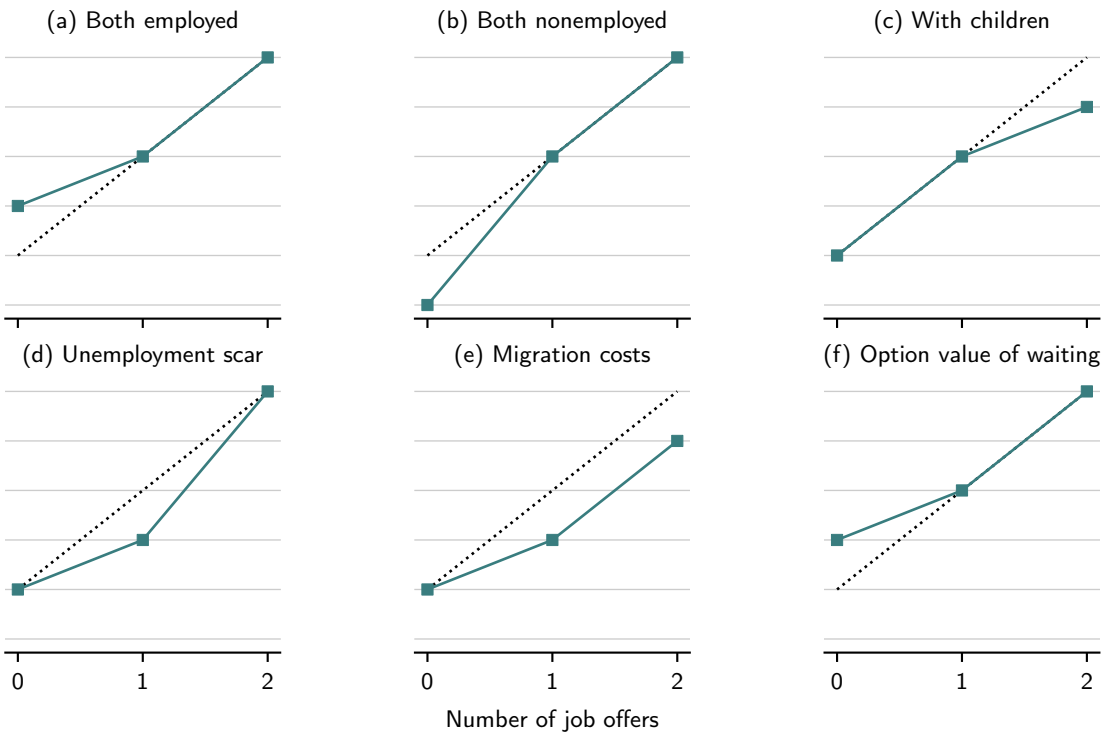

Figure 2: Illustration of economic forces shaping the curvature of $V - \chi$. The coordination friction is slack in the cases depicted in panels (b) and (c), and binds in the other cases. Solid lines depict $V - \chi$, while dashed lines depict a linear benchmark absent the respective force.

Finally, consider the case where the search elasticity across locations η is finite (Panel f). In this case, there is an option value of delaying migration reflected in $\Omega(N^{\frac{1}{1+\eta}}) > 0$. Intuitively, with $\eta < \infty$, there are efficiency gains from broadening search to multiple locations. After migrating, these efficiency gains are lost when the trailing spouse narrows their search to the location of the leading spouse.[13] Again, this introduces convexity and makes it more likely that the coordination friction binds.

4 Quantitative Model

We now introduce the quantitative model and calibrate it to the U.S. labor market. The quantitative model is a special case of the general framework in Section 2, with locations corresponding to U.S. commuting zones and households being differentiated by several characteristics beyond their employment status and residence.

[13] This is true regardless of whether repeated migration occurs along the equilibrium path. In either case, because search and migration costs are sunk, there is a positive option value of delaying migration if η is finite.

4.1 Setup

Household heterogeneity We quantify our model for dual-earner households whose adult members are composed of one woman, indexed by $i = f$, and one man, indexed by $i = m$. Households are differentiated by a vector $\boldsymbol{s} = (o_f, o_m, h_f, h_m, k)$, which along with households' employment status and their location defines their type $(\boldsymbol{e}, \boldsymbol{s}, r)$. Specifically, each spouse i is characterized by an immutable occupation, o_i, and a time-varying human capital level, $h_i \in \{\underline{h}, \overline{h}\}$.[14] In addition, we differentiate households with and without children, $k \in \{0, 1\}$.

Shocks Beyond search and migration, household types $(\boldsymbol{e}, \boldsymbol{s}, r)$ evolve stochastically through several independent shocks. The arrival rates of shocks are indexed by households' *occupation vector*, $\boldsymbol{o} \equiv (o_f, o_m)$, allowing them to vary across households.

First, jobs are destroyed at an exogenous, gender-specific rate $\delta_i(\boldsymbol{o})$. Second, children enter and exit households at rates $\pi_{k\uparrow}$ and $\pi_{k\downarrow}$, respectively. Third, human capital appreciates or depreciates as a function of a worker's current employment status: Employed workers' human capital appreciates to $h_i = \overline{h}$ at rate $\pi_{h\uparrow}(\boldsymbol{o})$, while non-employed workers' human capital depreciates to $h_i = \underline{h}$ at rate $\pi_{h\downarrow}(\boldsymbol{o})$. Finally, households are exposed to location preference shocks (detailed below) that induce them to relocate from commuting zone r to q at rate $\pi_{q|r}(\boldsymbol{o})$.

Preferences and technology The instantaneous utility and labor product are given by

$$u(\boldsymbol{e}, \boldsymbol{s}, r) = \sum_{i \in \{f,m\}} \underbrace{\left\{ h_i z_i(o_i, r) \cdot \mathbb{1}_{e_i=1} + b_i z_i(o_i, r) \cdot \mathbb{1}_{e_i=0} \right\}}_{\text{labor product + home production}} + \underbrace{a(k, r) - p(r)}_{\text{amenities − rent}} - \underbrace{\xi(r) \cdot \mathbb{1}_{k=e_f=e_m=1}}_{\text{childcare cost}}.$$

Here, $z_i(o_i, r)$ is a gender-, occupation- and commuting-zone-specific productivity that scales the labor and home products, $h_i z_i(o_i, r)$ and $b_i z_i(o_i, r)$. The value of living in commuting zone r is further determined by the value of its amenities $a(k, r)$, which differs by child status k, its cost of living $p(r)$, and its childcare costs $\xi(r)$. The latter accrue only if a household has children and both spouses are employed.

[14] Our model rules out occupational mobility, consistent with the broad occupational categories used in the calibration below; allowing for finer occupations or occupational switching would substantially expand the household state space, challenging computational feasibility. While occupational switching may contribute to the cost of becoming a trailing spouse, as emphasized by Rueda & Wilemme (2025), our calibration disciplines this cost directly through reduced-form wage-scar moments that capture the total earnings losses following job loss, including those arising from occupational switching.

4.2 Calibration

We calibrate the model using household data from the American Community Survey (ACS) and the Current Population Survey (CPS), which we combine with geolocation data, commuting zone data from the Opportunity Atlas (Chetty et al. 2025) and our own web-scraped data that inform the geography of the U.S. labor market. Appendix B.1 describes these data sources.

The calibration is organized around the two colocation constraints described above. Wage geography disciplines location-specific productivities and therefore geographic misalignment. CPS worker flows discipline search efficiency and job-separation rates. Wage scars discipline human-capital dynamics, which determine the cost of becoming a trailing spouse. Migration flows discipline migration costs. Finally, the cross-sectional distribution of households across commuting zones disciplines residual location-preference shocks.

Geography We consider all commuting zones (CZ) in the 48 contiguous states, excluding a few isolated commuting zones for which we lack the data to construct our amenity measure.

For computational efficiency, we merge commuting zones that are close geographically and in terms of all observational characteristics. We describe the details of our merging algorithm in Appendix B.4. Starting from 720 commuting zones with non-missing data, we obtain a final data set of 546 commuting zones post-merge.

Location preference shocks We design the arrival rates of location preference shocks, $\pi_{q|r}(\boldsymbol{o})$, to match the cross-sectional distribution of households by occupation vector $\boldsymbol{o}$ across commuting zones. Throughout the paper, we interpret location-preference shocks as residual migration caused by *non-work* forces, such as family ties and other idiosyncratic preferences. They are not included in our definition of *work-related* migration, which is the focus of the analysis below.

To invert the cross-sectional distribution in a computationally feasible way, we assume that location shocks induce a lump-sum utility shift,

$$\tau(\boldsymbol{e}, \boldsymbol{s}, r, q) = -\Big(V(\boldsymbol{e}, \boldsymbol{s}, q) - \chi(q|\boldsymbol{s}, r) - V(\boldsymbol{e}, \boldsymbol{s}, r)\Big),$$

that makes a household exactly indifferent to migrate from r to q.[15] With this design, households' value function and search policies, characterized by (4), are independent of the distribution of location shocks $\{\pi_{q|r}(\boldsymbol{o})\}$. For a given model parameterization, we can thus solve for V and the endogenous search and migration policies in a first step, and then

[15] To avoid indirect effects on the employment distribution, we assume that location shocks do not alter the employment status $\boldsymbol{e}$.

Table 1: Frequency distribution over occupation vectors

o_f/o_m	Managers/prof/tech	Clerical/retail sales	Low-skill services	Production/crafts	Machine operators	Construction/farm	Total
Managers/prof/tech	.32	.06	.03	.02	.02	.10	.55
Clerical/retail sales	.11	.04	.02	.01	.01	.09	.28
Low-skill services	.04	.02	.03	.01	.01	.06	.17
Total	.47	.12	.08	.04	.04	.25	1.00

Notes.—Frequencies are among dual-earner households in the ACS. Rows correspond to women's occupation and columns to men's occupation, classified following Autor & Dorn (2013).

design $\pi_{q|r}(\boldsymbol{o})$ in a second step to match, at the steady state, the empirical distribution over commuting zones.

We do so by choosing the distribution of shocks $\{\pi_{q|r}(\boldsymbol{o})\}$ with the smallest total prevalence that still allows us to exactly match the empirical distribution over commuting zones and occupation vectors. Our algorithm does so in a computationally efficient way: On an Intel i7-9700 computer, solving for V, identifying the set $\{\pi_{q|r}(\boldsymbol{o})\}$ that matches the empirical distribution over $\boldsymbol{o}$, and solving for the implied ergodic distribution over $(\boldsymbol{e}, \boldsymbol{s}, r)$ takes about 2.5 seconds.

Assigned parameters We parameterize the model at a monthly frequency. Households retire and are replaced by new households at a monthly rate of $0.021/12$, chosen so that the average household's work life lasts for 47 years. The effective discount rate ρ is set to the retirement rate plus $0.05/12$, corresponding to an annual time preference rate of 5%.

We categorize households using the occupation classification by Autor & Dorn (2013).[16] To economize on states, we drop occupations that are pursued by less than 3% of workers per gender.[17] This yields 18 occupation vectors, $(o_f, o_m) \in \{1, \dots, 3\} \times \{1, \dots, 6\}$, that cover 93.5% of all dual-earner households in the ACS. Tables 1 and 2 show their frequency distribution and average wages.

We set the job separation rates $\delta_i(\boldsymbol{o})$ to the empirical employment to non-employment rates

[16]Specifically, we differentiate occupations by the six top-level occupation classifications in Autor & Dorn (2013): (1) managers/professionals/technicians/finance/public safety, (2) clerical/retail sales, (3) low-skill services, (4) production/crafts, (5) machine operators/assemblers, (6) transportation/construction/mechanics/mining/farm. Henceforth, we reference these occupation categories using shortened labels.

[17]Following this criterion, we drop three out of six top-level occupations for women ("production/ crafts", "machine operators", and "construction/ farm") and do not drop any occupations for men.

Table 2: Average wage by occupation and gender

Occupation	Women	Men
Managers/prof/tech	1.57	2.45
Clerical/retail sales	.81	1.25
Low-skill services	.52	.70
Production/crafts	–	1.35
Machine operators	–	1.05
Construction/farm	–	1.12

Notes.—Wages are normalized relative to the median ACS wage. Occupations are classified following Autor & Dorn (2013).

in the CPS, measured at the gender × occupation-level.[18] We use a Cobb-Douglas matching function, $\lambda(\theta) = \theta^{\gamma}$, with matching elasticity $\gamma = 0.2$, as estimated by Lange & Papageorgiou (2020). We set the search elasticity across locations to $\eta = 0$, so search effort is inelastic across submarkets, as in Albrecht et al. (2006), Kircher (2009) and Galenianos & Kircher (2009). Thus, conditional on searching, households allocate equal search effort across targeted locations. Job-finding rates remain endogenous, however, because the value split determines the market tightness $\theta_{i,q}$ targeted in each submarket. At the extensive margin, households search only where the resulting search gains exceed migration costs. For computational efficiency, we limit workers to a maximum of four simultaneous searches and verify ex post that the restriction is locally non-binding almost everywhere.[19]

Next, we normalize the high human capital realization to $\overline{h} = 1$ and set $b_m = 0.4$, so the ratio between home and labor product equals 0.4 for high-h men (c.f., Shimer 2005). We then set women's home product $b_f = b_m \cdot \mathrm{E}[z_m]/\mathrm{E}[z_f] = 0.625$, so the average home product is the same across genders. With b pinned down, we set $\underline{h} = b_m = 0.4$, capturing the idea that low-h men are marginally attached to the labor market. Whether low-h men join the labor force then depends on the value gain from future human capital appreciation versus the childcare costs that incur if they have children and their spouse is employed.

Our calibration of local productivities, $z_i(o_i, r)$, exploits that for almost all (i, o_i, r) the median worker in our model has high human capital.[20] With this in mind, we infer the local productivities $z_i(o_i, r)$ from the gender × occupation-specific median wages in a commuting

[18] Because endogenous separations are quantitatively rare in the calibrated model, total model-implied separations closely approximate the targeted employment to non-employment rates.

[19] We verify this in the calibrated model by increasing the maximum number of simultaneous searches to five. We find that less than 0.001% of workers search in more than four markets simultaneously.

[20] Intuitively, the median worker is determined by the appreciation and depreciation rates $\pi_{h\uparrow}$ and $\pi_{h\downarrow}$, which we estimate below to match the *average* unemployment scar in the data. Given $\overline{h} - \underline{h}$, which is about five times the size of the average impact scar, our estimated process implies that the majority of workers has "high" (or, more accurately, "median") human capital. We verify numerically that conditional on (i, o_i, r), the median workers are indeed of type $\overline{h}$ for almost all gender × occupation × commuting zone combinations.

zone r, as measured in the ACS.[21] We set their scale to normalize economy-wide average earnings to 1.

Next, we use average rents for two bedroom apartments to inform the cost of living $p(r)$ for each commuting zone. We set childcare costs $\xi(r)$ to 8.5% of the median household income in each commuting zone, consistent with the average, age-weighted household expenditures on childcare documented by Guner et al. (2020). We set the monthly arrival rate of children to the fertility rate among childless households in the ACS, $\pi_{k\uparrow} = 0.075/12$. We then use the departure rate of children to match the share of households with children in the ACS, which is 56%, yielding $\pi_{k\downarrow} = (0.56^{-1} - 1) \cdot \pi_{k\uparrow}$ net of the retirement rate.[22]

Finally, we combine several existing and web-scraped data sources to inform the local amenity values $a(k, r)$. Our data include information on crime rates, various climate and weather categories, walkability scores, measures of beach access and quality, and various data on infrastructure, such as hospital quality and school quality measures. We assume that all school-related measures are valued by households with children, while all other amenities are valued by all households. Appendix B.3 describes the data in detail.

One challenge in calibrating the amenity value is that $a(k, r)$ enters $u(\boldsymbol{e}, \boldsymbol{s}, r)$ in *income-equivalent* units, requiring us to convert the various data into income-equivalent units. To do so, we assume that regional differences in income-equivalent amenities have the same passthrough rate on local rents as regional differences in wages, and that the passthrough of school quality scales with the local share of households with children. Given these assumptions, we can infer the amenity value from the following regression:

$$p_r = \beta_0 + \epsilon \cdot \left(\overline{w}_r + \beta_1' \boldsymbol{a}_r^{\text{all}} + \beta_2' \overline{k}_r \boldsymbol{a}_r^{\text{kids}} \right) + v_r, \tag{12}$$

where p_r are average rents in commuting zone r, $\overline{w}_r$ is the local median household income, $\overline{k}_r$ is the local share of households with children, and $\boldsymbol{a}_r^{\text{all}}$ and $\boldsymbol{a}_r^{\text{kids}}$ collect our various amenity data (applicable to all households and only for households with children, respectively). We estimate a passthrough rate ϵ of 0.161, which is consistent with there being significant mobility frictions (as our framework indeed delivers). Given our estimate for ϵ, we infer the local

[21] Equating productivities with wages is model-consistent for two separate reasons. First, for labor contracts to be self-enforcing in the absence of contractual commitments on the worker-side, workers must be paid their labor product at all times other than during the instant they are hired (see, e.g., Menzio & Shi 2011; Baley et al. 2022). Second, while our calibration uniquely pins down the vacancy cost c relative to search budgets $\bar{\kappa}_i(\boldsymbol{e}, \boldsymbol{s})^{(1-\gamma)/\gamma}$, their absolute value is indeterminate. Without loss of generality, we can thus consider the limit where both $c \to 0$ and $\bar{\kappa}_i(\boldsymbol{e}, \boldsymbol{s}) \to 0$, in which case workers are always paid their labor product under the unique equilibrium labor contract.

[22] To be consistent with our "perpetual youth" setting, we re-weight the ACS sample when computing the fertility rate and the share of households with children so that age is distributed geometrically in the sample.

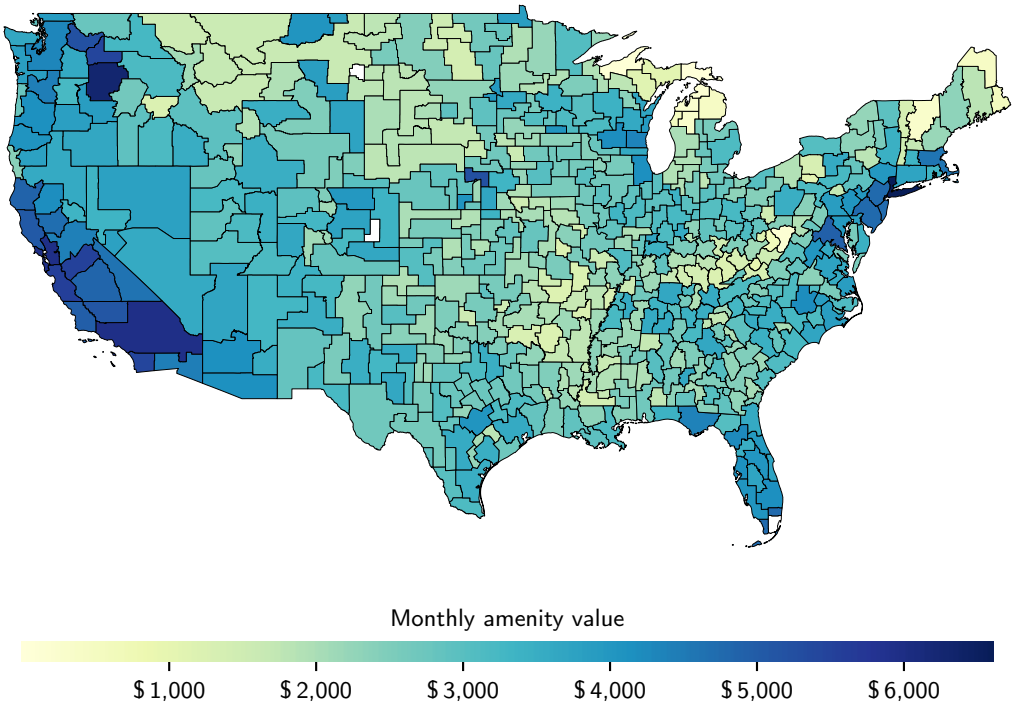

Figure 3: Income-equivalent amenity values for households without children. Amenity values are normalized to have a minimum of zero. White patches mark commuting zones with missing data.

amenity values as

$$a(k,r) = a_0 + \beta_1' \boldsymbol{a}_r^{\text{all}} + \beta_2' \boldsymbol{a}_r^{\text{kids}} \cdot \mathbb{1}_{k=1},$$

for some constant a_0. Without loss of generality, we normalize a_0 so that $\min_r a(0,r) = 0$.[23] Figure 3 plots the estimated amenity values for $k = 0$.

We summarize all exogenously assigned parameters in Table 3.

Estimated parameters We calibrate the remaining parameters using the method of moments, with weights chosen to minimize the relative distance between model and empirical moments. All model moments are computed at the steady state. Our estimation leverages that the steady state distribution is generated by 18 non-communicating Markov chains, separated by occupation vectors $\boldsymbol{o}$. In choosing target moments and parameters that are also indexed by $\boldsymbol{o}$, we can therefore split the estimation into 18 subproblems, greatly reducing the computational complexity. As usual, conditional on $\boldsymbol{o}$, all parameters are identified jointly. In the following, we provide a heuristic mapping from moments to parameters to guide intuition.

In our model, the strength of search frictions is determined by the search budgets relative to the vacancy cost, $\bar{\kappa}_i(\boldsymbol{e}, \boldsymbol{s})/c^{\gamma/(1-\gamma)}$, which are not separately identified. For some arbitrary

[23] Constant shifts in $a(k,r)$ translate to constant shifts in V by a_0/ρ and are thus of no consequence.

Table 3: Assigned parameters

Parameter	Value	Source
Time preference rate, annualized	.05	literature
Retirement rate, annualized	.021	avg. working life of 47 years
Home production, b_m, b_f	.4, .625	Shimer (2005), ACS
High human capital, $\overline{h}$	1.0	normalization
Low human capital, $\underline{h}$	.4	same as b_m, see text
Search elasticity across locations, η	.0	literature
Matching elasticity, γ	.2	Lange & Papageorgiou (2020)
Child arrival rate $\pi_{k\uparrow}$, annualized	.075	ACS
Child departure rate $\pi_{k\downarrow}$, annualized	.038	ACS
Job separation rates, $\{\delta_i(\boldsymbol{o})\}$	see text	CPS
Local productivities, $\{z_i(o_i, r)\}$	see text	ACS
Cost of living, $\{p(r)\}$	see text	Opportunity Atlas
Childcare costs, $\xi(r)$	see text	Guner et al. (2020)
Amenities, $\{a(k, r)\}$	see text	Opportunity Atlas, web scraped

normalization of c, we parameterize $\bar{\kappa}_i(\boldsymbol{e}, \boldsymbol{s}) = g_\kappa(i, \boldsymbol{o})$, and use it to match the empirical job-finding rate out of unemployment in the CPS, measured at the gender $\times$ occupation-pair level. To be consistent with the data, we only consider non-employed workers that are actively searching for jobs when computing the job-finding rate in the model.

Next, following Jung & Kuhn (2019) and Jarosch (2023), we estimate the human capital appreciation and depreciation rates, $\pi_{h\uparrow|e}(\boldsymbol{o})$ and $\pi_{h\downarrow|u}(\boldsymbol{o})$, to match the empirical steepness of job ladders. To do so, we simulate the wage scar of male workers separated from their job at $t = 0$ relative to the control group of non-separated workers, $\log(w_t^{\text{treat}}/w_t^{\text{control}})$, and match it to the estimated wage scars in Huckfeldt (2022) for $t = 12$ and $t = 36$ months.[24]

It remains to calibrate the migration costs $\chi(q|\boldsymbol{s}, r)$. To do so, we differentiate between "work-related" migration, which in the model corresponds to the endogenous migration through job search, and other "residual" migration, captured in the model through location shocks. On the empirical side, we consider 45% of all observed migration as work-related, based on the survey evidence in Maurer (2017).[25]

With this distinction, we calibrate migration costs by targeting, for each occupation vector $\boldsymbol{o}$, both the rate of work-related migration and its distribution in the ACS. Specifically, most migration in the data occurs at short distances and between commuting zones with similar population sizes. Motivated by these empirical patterns, we parameterize $\chi(q|\boldsymbol{s}, r)$ in terms of the spatial distance between any two commuting zones' population-weighted centroids,

[24]The wage scars documented in Huckfeldt (2022) may partly reflect endogenous responses to job loss, such as occupational switching. We load these losses onto human-capital depreciation because the relevant calibration target for our counterfactuals is the total equilibrium cost of job loss, not the particular mechanism through which that cost arises.

[25]This is closely in line with our own calculation based on a corresponding survey question in the CPS, which indicates that 47% of across-state migration is for a new job or a job transfer.

Table 4: Summary of target moments

Moment	Level	No. of moments	Source	$\approx$ maps to
Job-finding rate	$o_f \times o_m \times$ gender	36	CPS	$\bar{\kappa}_i(\boldsymbol{o})$
Wage scar, 1 & 3 yrs	$o_f \times o_m \times$ year	36	Huckfeldt (2022)	$\pi_{h\uparrow}(\boldsymbol{o})$, $\pi_{h\downarrow}(\boldsymbol{o})$
Migration rate, by bin	$o_f \times o_m \times$ bin	144	ACS	$\chi(q\|\boldsymbol{o}, r)$
Distribution over CZs	$o_f \times o_m \times r$	9828	ACS	$\pi_{q\|r}(\boldsymbol{o})$

Notes.—Due to adding up constraints on target distributions, the number of linearly independent moments is reduced by 36, matching the number of estimated parameters.

$d^{\text{geo}}(r, q)$, and the absolute difference in their population sizes, $d^{\text{pop}}(r, q)$.[26] Specifically, we set

$$\chi(q|\boldsymbol{o}, r) = \sum_{k\in\{\text{geo,pop}\}} \sum_{j\in\{1,\ldots,4\}} \chi_j^k(\boldsymbol{o}) \cdot \mathbb{1}_{d^k(r,q)\in \text{Bin}_j^k},$$

with four geographic and four population bins, $\{\text{Bin}_j^{\text{geo}}, \text{Bin}_j^{\text{pop}}\}_{j=1}^4$, chosen to capture the cross-bin variation in migration rates (cf. Figure 5). For each occupation vector $\boldsymbol{o}$, we normalize $\chi_1^{\text{geo}}(\boldsymbol{o}) = 0$ and then estimate the seven remaining cost parameters, $\{\chi_j^k(\boldsymbol{o})\}$, to match the total work-related migration rate of occupation vector $\boldsymbol{o}$ as well as its frequency distribution over the four geographic and the four population bins.

Table 4 summarizes the moments targeted in our estimation.

4.3 Model fit and validation

Targeted moments Our algorithm ensures that the model matches the household-by-occupation vector distribution across commuting zones by construction. Figures 4 and 5 compare the remaining 216 moments with their empirical targets. The model achieves a close fit for these moments as well.

Untargeted moments For independent validation, we compare the estimated model's predictions for employment rates and gender wage gaps with their empirical counterparts, both at the steady state and following migration events.

First, consider the steady-state moments (Table 5, Columns 2 and 3). As expected, the estimated model nearly replicates the gender wage gap, which is closely linked to the calibrated productivities.[27] More noteworthy is that the model also closely matches steady-state employment rates. These reflect labor force participation decisions shaped by the share

[26] Potential underlying mechanisms include state dependence in location preferences, whereby households develop attachment to locations of a given size, or persistent heterogeneity in preferences over urban versus rural environments.

[27] The wage gaps fall toward the upper end of those reported in the literature (see Doepke & Tertilt 2016 for an overview). This is because we include part-time workers in all wage calculations, both when calibrating productivities and when computing the empirical gaps.

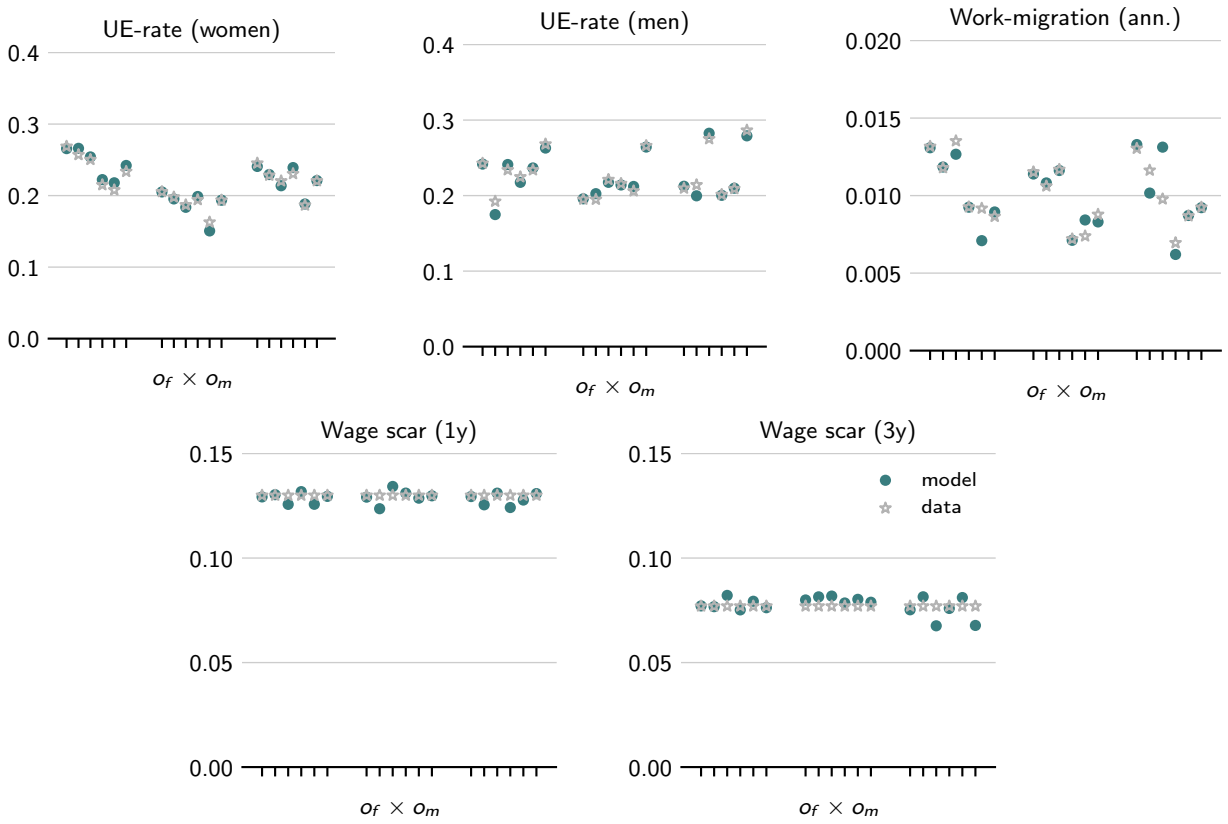

Figure 4: Targeted moments 1: Job-finding rates, annualized migration rates, and wage scars by occupation vectors. Occupation vectors are grouped by women's occupation, $o_f \in \{1, 2, 3\}$, with individual ticks corresponding to men's occupation, $o_m \in \{1, \ldots, 6\}$.

of low-h workers, their relative home versus market productivity, the option value of gaining work experience, and childcare costs — all calibrated independently of employment rates.

Next, consider the same moments after migration. Because trailing spouses often find jobs within weeks of moving, comparing gender gaps at the instant of migration ($t = 0$) to time-aggregated data would be misleading. To address this, we compare the model with data three months post migration (Table 5, last two columns). Again, employment rates and gender gaps align closely with their empirical analogs.

Offer storage and coordinated matching This comparison indirectly tests our assumption that the baseline model excludes offer storage and coordinated matching. If households commonly stored offers or employers routinely coordinated hiring across spouses, the baseline model would underpredict dual employment shortly after migration. However, the estimated model almost perfectly replicates the observed dual-employed share at three months (Table 5, bottom row), suggesting little scope for offer storage or coordinated matching in our setting.

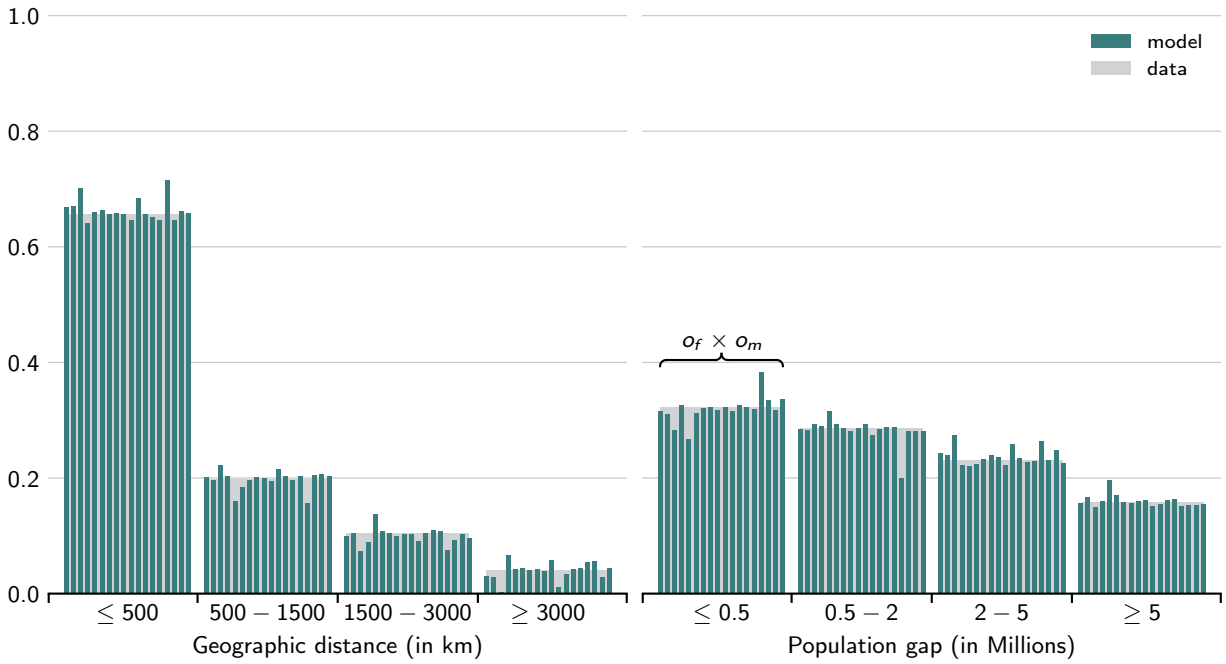

Figure 5: Targeted moments 2: Migration frequency over geographic distance and population gap.

5 Dual-Earner Migration and Gender Inequality

As shown in Table 5 above, gender inequality is substantial and becomes even more pronounced after migration. Measured in earnings, the steady-state gap is 46%, which widens to 80% immediately following migration — largely because in four out of five migration events, women are the trailing spouse. As a result, the model implies that the three-year present value gain from migration is substantially larger for men than for women: 0.13 versus 0.02 (in units of

Table 5: Untargeted moments: Employment rates and gender gaps

	Steady state		Post Migration ($t = 0$)		Post Migration ($t = 3$)	
Moment	Model	Data	Model	Data	Model	Data
Earnings gap	.46	.45	.80	n/a	.60	.68
Wage gap	.31	.34	.30	n/a	.41	.49
Employment rate						
women	.73	.73	.22	n/a	.60	.59
men	.93	.88	.78	n/a	.87	.91
Dual-employed share	.68	.65	.00	n/a	.48	.52

Notes.—Gender gaps are defined as $1 - \mathrm{E}[x_f]/\mathrm{E}[x_m]$. Empirical moments are computed in the PSID, unconditionally and 3 months after a work-related migration event has been reported.

average lifetime earnings).[28] We next explore how pre-existing gender differences in model inputs translate into these gaps.

Decomposing gender inequality In our calibration, three inputs are gender-specific: productivities $z_i(\cdot)$, job separation rates $\delta_i(\cdot)$, and job-finding rates. The first two enter the model directly, while the third informs the estimation of search capacities $\kappa_i(\cdot)$. In the absence of gender differences in z_i, δ_i, κ_i, the model would not generate any gender gaps.

To assess the contribution of each input to the estimated gender gaps, we re-simulate the model, equalizing these parameters across genders in all possible permutations. We further break down the role of productivity differences by separately equalizing the average productivity level and the dispersion of productivities across commuting zones, giving us four gender-specific inputs in total. Our decomposition is exact and is independent of the order of variation.

Formally, denote a permutation of the four gender-specific inputs by

$$\boldsymbol{x} = (\mathbb{1}_k)_{k\in\{1,\dots,4\}} \in \Big\{(0,0,0,0),(0,0,0,1),\dots,(1,1,1,1)\Big\},$$

where $\mathbb{1}_k$ indicates that input k is equalized across genders by setting the corresponding parameter for women to the one for men. Given a permutation $\boldsymbol{x}$ with $\boldsymbol{x}_k = 0$, the marginal contribution of input k to gap $\mathcal{G}$ is given by

$$\mathrm{d}_k\mathcal{G}(\boldsymbol{x}) = \mathcal{G}(\boldsymbol{x} + \mathbb{1}_k) - \mathcal{G}(\boldsymbol{x}).$$

Our decomposition averages the marginal contribution for each input across permutations. To ensure that the decomposition adds up exactly to the total contribution, we weigh each of these permutations by the number of paths connecting from $(0,0,0,0)$ to $(1,1,1,1)$ through $(\boldsymbol{x}) \to (\boldsymbol{x} + \mathbb{1}_k)$. That is,

$$\mathit{contrib}_k = \sum_{\boldsymbol{x}\in\{\boldsymbol{x}:\boldsymbol{x}_k=0\}} w(\boldsymbol{x}) \cdot \mathrm{d}_k\mathcal{G}(\boldsymbol{x})$$

with

$$w(\boldsymbol{x}) = \frac{\overbrace{\Big(\sum_l \boldsymbol{x}_l\Big)!}^{\text{paths in}} \times \overbrace{\Big(n - \sum_l \boldsymbol{x}_l - 1\Big)!}^{\text{paths out}}}{\underbrace{n!}_{\text{all paths}}}, \qquad n = 4.$$

[28] We compute these gains by simulating a representative sample of migrating households and comparing it with the same sample absent migration.

Table 6: Decomposition of gender gaps

	Steady state		Post Migration		
	Employment gap	Wage gap	Employment gap	Wage gap	Leading spouse gap
Productivities (level)	17 %	28 %	40 %	17 %	17 %
Productivities (dispersion across CZs)	1 %	1 %	9 %	8 %	3 %
Job-finding rates	−4 %	−1 %	9 %	4 %	2 %
Job-separation rates	8 %	3 %	15 %	1 %	5 %
Total	21 %	31 %	72 %	30 %	28 %

Notes.—Post-migration gaps are computed immediately after migration ($t = 0$). Employment and wage gaps are defined as $1 - \mathrm{E}[x_f]/\mathrm{E}[x_m]$. The leading spouse gap is defined as the share of work-related relocations in which the man is the leading spouse, in excess of 50%.

Table 6 summarizes the average marginal contributions, $contrib_k$. Level differences in productivity explain most of the gender gaps. Equalizing average productivities, men would lead 60% (instead of 78%) of moves. Larger job separation rates for women account for most of the remaining gap, while job-finding rates and productivity dispersion matter mainly for the post-migration gaps.

The productivity component should be interpreted broadly. In the model, gender-specific productivities capture any force that generates systematic wage differences across spouses conditional on the state variables we track, including discrimination, differences in bargaining power, compensating differentials, and hours choices. In this decomposition, the component also absorbs wage differences associated with the broad occupational categories used in the calibration.

Is the main earner always the leading spouse? Given that wage levels explain most of the leading-spouse gap, it is natural to ask whether the main earner, defined as the spouse with the higher labor product, $h_i z_i(o_i, r)$, typically leads the move. Table 7 shows that this is true for main-earning men, but not for main-earning women, who trail their spouses in about half of all moves.

This pattern aligns with evidence from Jayachandran et al. (2024), who find that cross–CZ moves improve men's labor market outcomes more than women's, regardless of which spouse has the greater earnings potential. In our model, this arises because men face larger productivity differences across CZs and accumulate more human capital post-migration, amplifying their returns to relocating.

Thus, the leading-spouse pattern is not a mechanical consequence of current earnings. Instead, it reflects how the two-body problem interacts with spatial returns, human-capital dynamics, and post-migration search.

Table 7: Leading spouse by main earner status

	Man is leading spouse	Woman is leading spouse
Man is main earner	92.0 %	8.0 %
Woman is main earner	45.6 %	54.4 %

Notes.—Shares of migration events led by men and women, conditional on who is the main earner. Rows sum to 100%.

Table 8: Leading spouse by occupation vector

o_f/o_m	Managers/prof/tech	Clerical/retail sales	Low-skill services	Production/crafts	Machine operators	Construction/farm
Managers/prof/tech	87 %	67 %	27 %	65 %	48 %	49 %
Clerical/retail sales	97 %	85 %	72 %	87 %	80 %	80 %
Low-skill services	98 %	90 %	81 %	94 %	87 %	83 %

Notes.—Share of migration events led by men. Rows correspond to women's occupation and columns to men's occupation, classified following Autor & Dorn (2013).

Heterogeneity by occupation Another natural question is whether leading-spouse gaps persist across occupations. Table 8 indicates that they largely do: Men lead most moves regardless of their own or their spouse's occupation. Women lead a substantial share only when they hold a high-paying occupation — such as in management, technology, or finance — while their spouse works in a low-paying occupation, such as low-skill services, construction or farming.

The role of children Finally, we examine how children affect migration. Table 9 shows that households with children migrate less often than childless households. In our framework, this reflects two forces: First, younger, childless households are less well matched and therefore have stronger incentives to relocate. Second, the arrival of children shifts amenity preferences and lowers the opportunity cost of having a trailing spouse, who can provide childcare.[29] The second force increases migration among households with children, but is dominated by the first force, implying an overall higher migration rate among childless households. This is consistent with evidence from the ACS, where childless households migrate at roughly twice the rate of households with children, though the model implies a somewhat steeper child gradient than in the data.

[29]While we do not explicitly model age, households enter our model without children, so that average time in the model does indeed correlate with child status.

Table 9: Work-related migration with and without children

	Model	Data
Children	.35 %	1.05 %
No children	2.09 %	2.26 %

Notes.—Annualized migration rates by child status. Data are cross-CZ migration rates in the ACS, scaled by the CPS share of migration due to new jobs or job transfers.

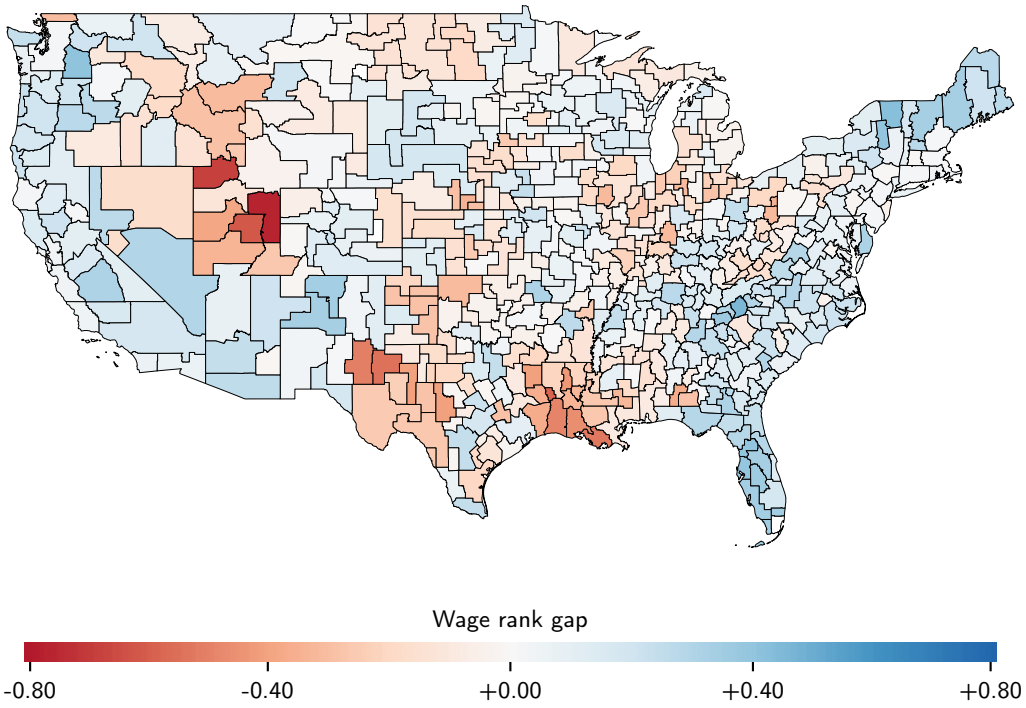

Figure 6: Empirical differences in wage ranks, $rank_{f,o}(r) - rank_{m,o}(r)$, averaged across occupation pairs. Each rank denotes the percentile rank of a commuting zone in the gender- and occupation-specific productivity distribution across commuting zones. Ranks are scaled to lie between 0 and 1. Data source: ACS.

6 Quantifying Colocation Constraints

Having shown that gender inequality widens after migration, we now quantify the extent to which colocation constraints contribute to this widening, and further affect household mobility, long-run sorting across locations, and earnings more broadly.

6.1 Geographic Misalignment

Figure 6 shows that misalignment is pronounced in the U.S. labor market, with gaps in spouses' earnings-prospects ranks reaching up to 80 percentile ranks in some commuting zones. Broadly, California, the Northeast, and Florida tend to rank higher for women, whereas the

Table 10: Overall losses attributed to geographic misalignment

	Women	Men
Earnings	1.22 %	.45 %
Employment	.00 %	−.02 %
Local productivity	.98 %	.68 %
Human capital	.00 %	−.07 %

Notes.—Losses are computed by comparing average earnings in the baseline economy with the geographically aligned counterfactual economy. Both economies are evaluated at their respective steady-state distribution.

Midwest and Texas tend to rank higher for men. Averaging across occupation pairs, the correlation between female and male productivity ranks is 0.56.

To quantify the role of geographic misalignment, we compare the baseline economy with the geographic alignment benchmark introduced in Section 3.1. Specifically, for each occupation pair, we hold the geographic distribution of male productivities fixed and reshuffle female productivities across commuting zones so that female productivity ranks coincide with the corresponding male ranks. All other parameters are held fixed at their baseline values; we recompute equilibrium search policies and value functions, and evaluate outcomes at the resulting counterfactual steady-state distribution.

Removing geographic misalignment raises earnings considerably, with larger effects for women than for men, as shown in Table 10. These earnings gains arise primarily from households sorting into higher-productivity locations. Employment and human capital remain essentially unchanged for women and decline slightly for men, playing little role in the earnings response. Despite the sizable earnings gains, the effect on mobility is modest, with annual migration rates rising from 1.12% to 1.18%. This is because, at the steady state, households are sorted into higher-productivity locations, partly depleting the higher gains from migration.

Overall, geographic misalignment lowers earnings and amplifies gender inequality by distorting where dual-earner households choose to live.

6.2 Coordination Friction

We next turn to the coordination friction. Unlike geographic misalignment, which changes the value of destinations, the coordination friction changes the risk associated with moving as illustrated in Section 3.2. At the steady state distribution, the coordination friction binds in nearly all non-local searches ($q \neq r$) and is slack in most local searches ($q = r$). We next explore how this affects labor markets, beginning with the direct impact on migrating households.

Table 11: Consequences of migration for earnings

	ΔPV in 3-year earnings		ΔPV in lifetime earnings	
	Baseline economy	Correlated matching	Baseline economy	Correlated matching
Women	.021	.084	.296	.366
Men	.129	.148	.610	.632

Notes.—ΔPV denotes difference in the present value of earnings between a representative sample of migrating households and an identical sample without migration, denominated in units of economy-wide average lifetime earnings.

Direct Effect on Migrating Households To provide context, we first quantify returns to migration in the baseline economy. To do so, we simulate earnings gains for a representative sample of migrating households and compare them to an identical sample absent migration. Table 11 shows the difference in the present value of earnings between these two groups.

Over the first three years after migration, men's present value gain equals 0.129 (in units of average lifetime earnings), compared to only 0.021 for women. In the longer run, the gap narrows but remains substantial: 0.610 for men versus 0.296 for women. These gaps largely reflect the leading-spouse gap discussed earlier, amplified by slowing human capital accumulation for the trailing spouse.

To isolate the causal effect of the coordination friction, we next consider the same migrating sample of households, but now simulate hypothetical present value gains when the coordination friction is relaxed by allowing for joint job offers as implied by the correlated matching benchmark.[30] As expected, relaxing the coordination friction barely changes men's gains, but dramatically increases women's short-term earnings.

The direct cost of the coordination friction is therefore borne disproportionately by women, not because the friction is gender-specific, but because gender differences in earnings prospects and search dynamics make women more likely to be the trailing spouse. Overall, the friction reduces women's migration gains by 0.070 (again, in units of average lifetime earnings), with most losses (0.063) accruing in the first three years. In relative terms, this reduces women's long-term migration gains by 19% and accounts for 18% of the corresponding gender gap.

Discouraged migration and power couples Beyond shaping post-migration outcomes, the coordination friction may also prevent households from moving altogether, especially when both spouses have valuable jobs to preserve. To quantify its impact on migration rates, we relax the coordination friction for a representative sample of households and their descendants,

[30]Specifically, here we keep both the migrating sample of households and their migration destinations fixed, changing only the employment status immediately following migration according to the hypothetical correlation policies $\omega_k(\boldsymbol{e}, \boldsymbol{s}, r, V)$ implied by the correlated matching benchmark with V equal to its steady-state value in the baseline economy.

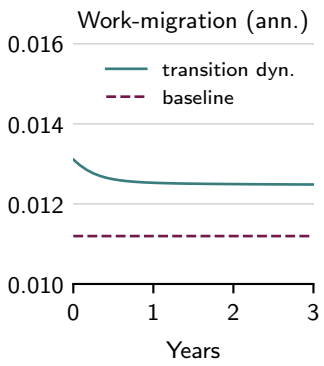

Figure 7: Impact of the coordination friction on work-related migration. The dashed line shows the annual steady-state migration rate in the estimated model. The solid line shows the rate for a representative sample of households for whom the friction is relaxed from $t = 0$ onward.

and study the resulting migration behavior. Formally, we compute the transition path from the baseline economy to the correlated matching benchmark, holding all primitives fixed at their baseline values. The transition is initialized at the baseline steady-state distribution, and search policies and value functions are updated along the path.

Figure 7 shows that relaxing the friction raises migration rates by 17% initially and by 13% in the long run, indicating a sizable effect on aggregate mobility.[31]

To understand the determinants of this increase in mobility, we next examine which households are discouraged from migrating when the friction is active. To do so, Table 12 contrasts households that migrate under the baseline with those that only migrate when the friction is removed. In line with the stylized model in Section 3.2, discouraged households tend to be dual-employed, high-income, and high-human-capital, resembling the stereotypical "power couple" (Costa & Kahn 2000). By contrast, baseline migrants are more likely to be childless and lower-income, making the friction less consequential for them.

Regarding the returns to migration, both groups have below-average local productivities, especially for men. Nevertheless, because human capital and productivity are complementary, discouraged households have significantly higher potential returns from migration. Consequently, earnings losses from discouraged migration are substantial: 0.140 for women and 0.168 for men (in units of average lifetime earnings). It is worth noting that while women bear most direct costs of the coordination friction, indirect effects from forgone migration opportunities affect both genders roughly equally.

[31]Note that by relaxing the friction perpetually, momentary migration incentives are reduced by the option value of migrating without friction in the future. Alternatively, we can measure the "compound effect" of the coordination friction on momentary migration rates by relaxing the friction during only a short time window and then reimpose it after. In this case, migration increases by over 60% to an annual rate of 1.86% while the friction is relaxed.

Table 12: Comparison of baseline migrants to discouraged migrants

	Migrating	Discouraged	Population average
Children	.14	.31	.56
Both employed	.38	.92	.68
Both non-employed	.22	.00	.02
Household income	1.21	2.20	2.00
Employment rate			
women	.52	.96	.73
men	.64	.96	.93
Human capital			
women	.77	1.00	.75
men	.89	.99	.87
Local productivity			
women	.92	.99	1.00
men	1.31	1.32	1.54

Notes.—Displayed are averages for three household groups: (i) Migrating: households who migrate in the baseline economy. (ii) Discouraged: households who do not migrate in the baseline economy, but migrate when the coordination friction is relaxed. (iii) Population average: all households, weighted by the steady-state distribution.

Combined effect To summarize, we quantify the overall impact of the coordination friction, taking into account the direct effect on post-migration employment, discouraged migration, as well as location choices. To do so, consider again the sample of households for whom we have relaxed the coordination friction at $t = 0$. Figure 8 plots transition dynamics in their earnings, employment, local productivity, and human capital. Relaxing the coordination friction immediately boosts women's employment and earnings, with further gains accruing through human capital accumulation and relocation to higher-productivity regions. By contrast, the combined effect of the coordination friction on male earnings is more subdued.

Overall, comparing steady states, the friction reduces women's steady-state earnings by

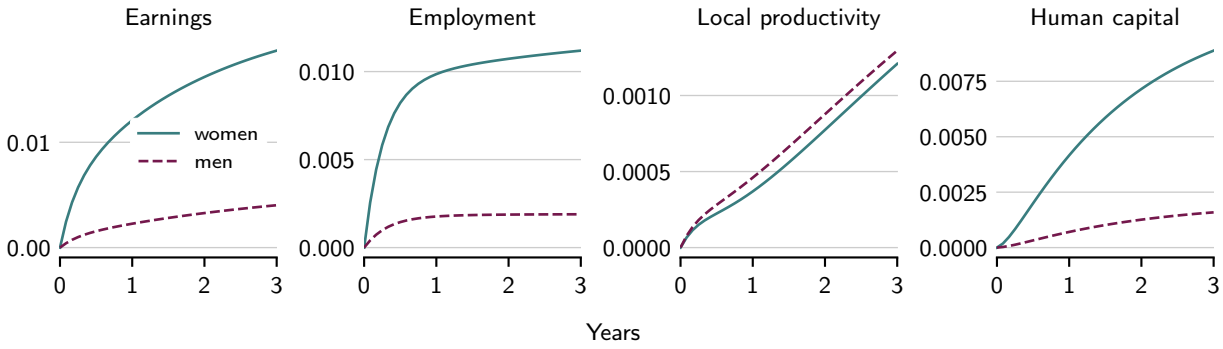

Figure 8: Impact of the coordination friction on earnings, employment, local productivities, and human capital. The plots show log-differences between a representative sample of households for whom the friction is relaxed from $t = 0$ onward and the steady state where the friction is active.

Table 13: Overall losses attributed to coordination friction

	Women	Men
Earnings	2.70 %	.86 %
Employment	1.29 %	.19 %
Local productivity	.54 %	.53 %
Human capital	1.18 %	.20 %

Notes.—Losses are computed by comparing average earnings in the baseline economy with the correlated matching benchmark. Both economies are evaluated at their respective steady-state distribution.

2.7% and men's by 0.9%, and further lowers amenity values equivalent to an additional 1% earnings loss. Table 13 decomposes the earnings losses into employment (1.3% for women, 0.2% for men), productivity losses from forgone relocation (0.5% for both), and human capital losses (1.2% for women, 0.2% for men).

7 Remote Work

We now study remote work as a relaxation of the two-body problem. By making jobs portable across locations, remote work enables workers to retain their jobs when the household moves. It therefore directly mitigates the coordination friction.

7.1 Model Extension

We incorporate remote work into the quantitative model described in Section 4.1. Apart from the modifications detailed below, the model remains unchanged.

Modeling remote work We model remote work as a job-level attribute that can evolve over the course of a match. Specifically, each job begins as non-remote and becomes remote at rate $\phi_i(\boldsymbol{o})$. Once a job becomes remote, it remains so for the duration of the match.

The defining feature of remote jobs is their portability across locations. If one spouse holds a remote job, they remain employed even if the household relocates to another commuting zone. Remote work thus directly relaxes the coordination friction, removing the need for trailing spouses to secure employment in the new location.

Mainly for tractability, we assume that remote workers' productivity is indexed to their current residence rather than to their residence at match-formation.[32]

[32]The assumption is tractable as it allows us to track each household's remote status using only two indicators $remote_f, remote_m \in \{0,1\}^2$, rather than tracking the origin or some "headquarter" location of jobs. From a computational perspective, the latter is challenging in a dual-earner setting, as it would expand the model's state space by a factor of 546^2 to about 94 Billion states. While the empirical determinants of remote workers' productivity remain the subject of ongoing debate (Bloom et al. 2015; Gibbs et al. 2023; Emanuel & Harrington 2024; Barrero et al. 2023), evidence by Brinatti et al. (2021) suggests that remote

Table 14: Share of remote workable jobs by occupation

Occupation	% potentially remote
Managers/prof/tech	57 %
Clerical/retail sales	55 %
Low-skill services	15 %
Production/crafts	–
Machine operators	–
Construction/farm	2 %
Women	49 %
Men	35 %

Notes.—Shares are computed by aggregating the occupation-specific shares of remote-workable jobs reported by Dingel & Neiman (2020) to the occupation classification by Autor & Dorn (2013). No remote workable jobs are reported for production/crafts and machine operators. The bottom panel aggregates these remote shares using the occupation–gender distribution in Table 1.

Calibration of job conversion rates The additional parameters that require calibration are the rates $\{\phi_i(\boldsymbol{o})\}$, at which non-remote jobs become remote. We calibrate these rates to match the occupation-level shares of jobs that can be performed entirely at home, as reported by Dingel & Neiman (2020). To align with our model's occupation structure, we aggregate these shares to the coarser occupational classification in Autor & Dorn (2013). Table 14 presents the aggregated shares.

We map these shares into conversion rates $\{\phi_i(\boldsymbol{o})\}$ by assuming that the share of remote jobs in each occupation is in steady state. That is,

$$\text{remote share}(o_i) = \frac{\phi_i(\boldsymbol{o})}{\phi_i(\boldsymbol{o}) + \delta_i(\boldsymbol{o})}, \tag{13}$$

where δ_i is the calibrated job separation rate.[33] Inverting equation (13) yields $\phi_i(\boldsymbol{o})$.

7.2 Implications of Remote Work

We now simulate the introduction of remote jobs. We initialize the economy at the steady state of the baseline without remote work. From $t = 0$ onward, jobs begin converting to remote jobs at the calibrated conversion rates $\{\phi_i(\boldsymbol{o})\}$. In the following, we explore the impact on migration and earnings along the transition path and at the new steady state.

Migration rates Remote work affects both the optimal timing and the long-run incidence of migration. Figure 9a plots migration rates along the transition path, following its introduction at $t = 0$. Migration initially drops by 28%, caused by the option value of waiting for a

wages indeed adjust with the residence of workers, reassuring our approach.

[33] In principle, equation (13) would also need to account for endogenous separations. However, since remote jobs survive migration, endogenous separations are negligible compared to δ_i.

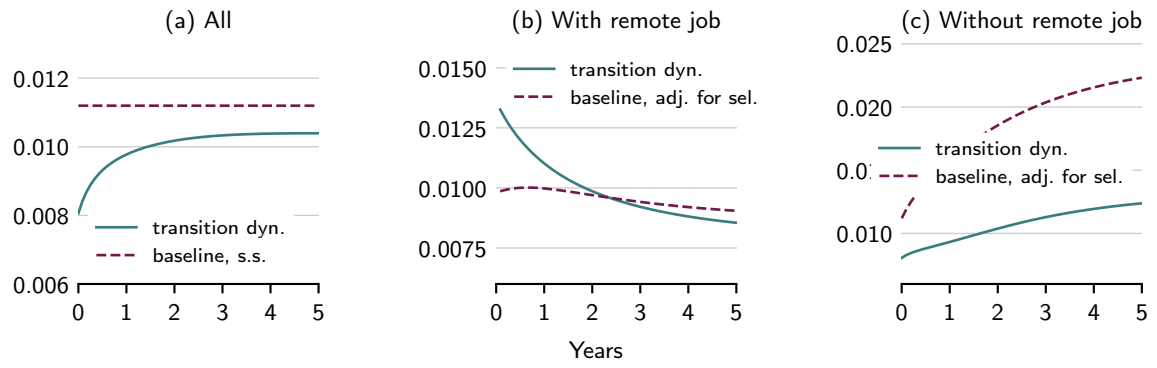

Figure 9: Transition dynamics following the arrival of remote work at $t = 0$. Panel (a) shows the annualized work-related migration rate along the transition path. Panels (b) and (c) break up the migration rate by remote status. Dashed lines correspond to migration in the baseline economy, computed at the steady state (panel a) and controlling for selection in $(\boldsymbol{e}, \boldsymbol{s}, r)$ conditional on remote status and time t (panels b and c).

remote job before relocating. As the incidence of remote jobs rises, so does the migration rate. Specifically, among households with at least one remote job, migration initially surges by 38% relative to their counterparts in the baseline economy (Figure 9b).[34] Over time, as these households relocate to preferred commuting zones, their migration rate declines, closing the gap to their baseline counterparts.

Households without remote jobs exhibit the opposite trend (Figure 9c). Caused by the option value of waiting, their migration rates initially fall, leading to increased geographic mismatch. This mismatch raises the potential gains from migration, which increases migration rates (solid line) as well as the option value of waiting (gap between solid and dashed line). At the remote economy's steady state, the option value reduces migration among non-remote households by 45%.

Taken together, aggregate migration converges to 7% below the baseline steady state, accounted for by the option value of waiting, partially offset by increased migration gains.

Composition of migrating households Table 15 further dissects the composition of migrating households in the remote economy's steady state. For comparison, columns 2 and 3 replicate the corresponding statistics from the baseline economy (cf. Table 12). Migrants with at least one remote job are typically dual-employed, high-income, and high-human-capital households — closely resembling the power couples who, in the baseline economy, were discouraged from migrating due to the coordination friction. In contrast, migrants without remote jobs exhibit average human capital and below-average employment levels, similar to the typical migrants in the baseline economy. Thus, the introduction of remote work shifts

[34] Formally, the comparison controls for selection in $(\boldsymbol{e}, \boldsymbol{s}, r)$ conditional on at least one spouse having a remote job, by computing migration in the baseline economy given the same conditional distribution over $(\boldsymbol{e}, \boldsymbol{s}, r)$ at each point in the transition path.

Table 15: Composition of migrants: baseline vs. remote work scenario

	Baseline economy		Remote work economy		
	Migrating	Discouraged	All migrating	Migrating w/ remote job	Migrating w/o remote job
Children	.14	.31	.16	.26	.10
Both employed	.38	.92	.46	.86	.20
Both non-employed	.22	.00	.21	.00	.35
Household income	1.21	2.20	1.34	2.16	.81
Employment rate					
women	.52	.96	.59	.93	.37
men	.64	.96	.66	.93	.48
Human capital					
women	.77	1.00	.83	.92	.77
men	.89	.99	.90	.95	.87
Remote job share					
women	.00	.00	.48	.77	.00
men	.00	.00	.28	.50	.00

Notes.—Displayed are averages for households who migrate in the baseline economy and in the remote work scenario. The compositions are computed at the respective steady states.

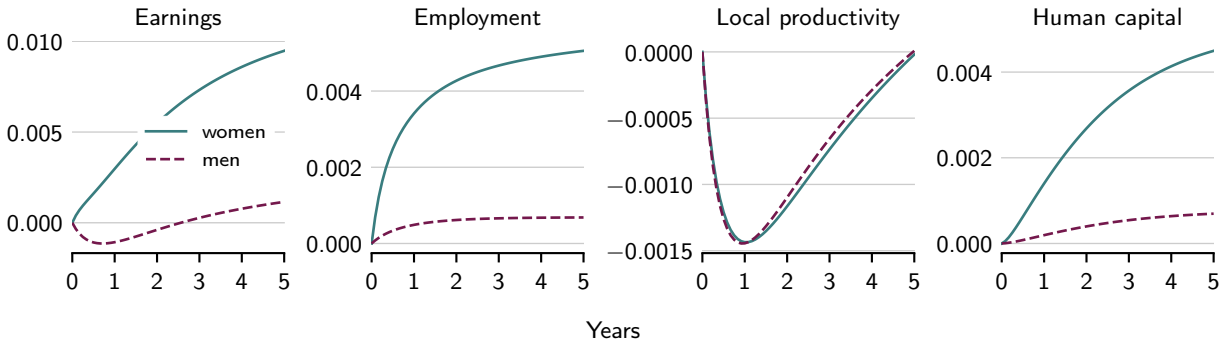

Figure 10: Impact of remote work on earnings, employment, local productivities, and human capital. The plot shows log-differences between a representative sample of households for whom positive transition rates to remote jobs are introduced from $t = 0$ onward and the steady state without remote work.

the migrant composition away from the baseline sample toward power couples with higher migration gains.

Earnings Figure 10 plots earnings and its components along the transition path. Overall, earnings trend up in the long run, driven by three factors. First, the ability to relocate without job loss for one's spouse directly increases employment and earnings, particularly for women, who are more likely to be trailing spouses at baseline. Second, the shift in migration toward power couples boosts earnings for both genders, due to complementarities between human capital and local productivity. Third, opposing these effects, postponed migration among non-remote households results in forgone earnings gains, explaining the initial decline

Table 16: Steady state comparison: baseline economy vs. remote work scenario

	Calibrated remote		Fully remote	
	Women	Men	Women	Men
Earnings	1.35 %	.37 %	2.36 %	.79 %
Employment	.55 %	.07 %	1.08 %	.14 %
Local productivity	.26 %	.24 %	.54 %	.54 %
Human capital	.52 %	.08 %	.95 %	.13 %

Notes.—Differences are computed between the steady-state distribution of the remote work economy relative to the baseline economy. "Calibrated remote" refers to the remote economy as calibrated in Section 7.1, "fully remote" refers to the case where all jobs are remote.

in productivity, reflected notably in male earnings.

Table 16 summarizes the steady-state impact on earnings, relative to the baseline. On average, earnings in the remote economy rise by 1.35% for women and 0.37% for men.

Comparison with losses from coordination friction From a modeling perspective, our implementation of remote work is closely related to the coordination friction: Both counterfactuals prevent employment losses among trailing spouses, without directly affecting wages. However, there are a few subtle differences.

First, unlike the correlated matching counterfactual — which reallocates job offers across households without altering the aggregate matching frontier — remote work expands the feasible set of matches.[35] Second, remote work relaxes the coordination friction only for employed workers, whereas the correlated matching counterfactual also treats unemployed workers.[36] Third, and most importantly, because remote work is limited to specific occupations, its aggregate impact is reduced due to a smaller treatment scope.

To control for differences in treatment scope, we also simulate a "fully remote" economy, in which all jobs are remote, as reported in the final two columns of Table 16. When all jobs are remote, the economy recovers most of the gains from eliminating the coordination friction (cf., Table 13). The remaining 10% gap reflects the continued impact of the friction on unemployed workers, partially offset by the expanded matching frontier of remote work.

In conclusion, remote work has the potential to eliminate the coordination friction for households most exposed to it, largely replicating the gains from directly removing it. In practice, accounting for occupational constraints on remote feasibility, remote work can cut average earnings losses attributed to the coordination friction by up to 50%.

[35]To see this, note that migration with a remote job is isomorphic to creating a new match without additional search effort by the worker and without creating a new vacancy.

[36]It is straightforward to show that the coordination friction is slack for all searches if at least one spouse has a remote job.

8 Concluding Remarks

This paper develops a spatial framework to study how search frictions shape job search, migration, and labor market outcomes of dual-earner households. We estimate the model for the U.S. labor market. The estimated model reproduces observed gender gaps in wages and employment without targeting them during the estimation.

We highlight geographic misalignment and a coordination friction as key constraints distinguishing dual-earner households from single-earner households. These constraints amplify gender inequality after migration, reduce migration rates — particularly among "power couples" — and lower steady-state earnings for both genders. We also examine how the rise of remote work mitigates these constraints, with the potential to offset a substantial share of these earnings losses in the long run.

Our framework is among the first to integrate dual-earner job search into a spatial labor market model. It is distinguished from the existing literature by its analytical tractability, which enables a large-scale estimation at the commuting-zone level. The estimated framework delivers rich predictions about dual-earner households' careers and gender disparities in employment and earnings.

The broader lesson is that household location constraints can generate labor-market misallocation even when each spouse searches optimally and households coordinate internally. Fundamentally, the misallocation arises because households are typically the unit of mobility, while jobs are matched to individual workers.

References

Albrecht, James, Pieter A. Gautier & Susan Vroman (2006). "Equilibrium Directed Search with Multiple Applications". *Review of Economic Studies* 73 (4), pp. 869–891.

Alon, Titan, Matthias Doepke, Jane Olmstead-Rumsey & Michèle Tertilt (2020). "This time it's different: the role of women's employment in a pandemic recession". NBER Working Paper (27660).

Alonzo, Davide, Nezih Guner & Claudio Luccioletti (2023). "Segregation and Sorting of U.S. Households: Who Marries Whom and Where?"

Autor, David H. & David Dorn (2013). "The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market". *American Economic Review* 103 (5), pp. 1553–97.

Azkarate-Askasua, Miren, Christian Hellwig & Miguel Zerecero (2023). "Home Bias and the Gender Wage Gap".

Bagga, Sadhika, Lukas Friedrich Mann, Ayşegül Şahin & Giovanni L. Violante (2025). "Job Amenity Shocks and Labor Reallocation". NBER Working Paper (33787).

Bagger, Jesper & Rasmus Lentz (2019). "An empirical model of wage dispersion with sorting". *The Review of Economic Studies* 86 (1), pp. 153–190.

Baley, Isaac, Ana Figueiredo & Robert Ulbricht (2022). "Mismatch Cycles". *Journal of Political Economy* 130 (11), pp. 2943–2984.

Barrero, José María, Nicholas Bloom & Steven J. Davis (2023). "The Evolution of Work from Home". *Journal of Economic Perspectives* 37 (4), pp. 23–50.

Berresheim, Ursula (2025). "Work from Home, Work for Less? How Workspace Flexibility Affects Mothers' Careers".

Blau, Francine D. & Lawrence M. Kahn (2017). "The Gender Wage Gap: Extent, Trends, and Explanations". *Journal of Economic Literature* 55 (3), pp. 789–865.

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

Braun, Christine, Charlie Nusbaum & Peter Rupert (2021). "Labor market dynamics and the migration behavior of married couples". *Review of Economic Dynamics* 42, pp. 239–263.

Brinatti, Agostina, Alberto Cavallo, Javier Cravino & Andres Drenik (2021). "The International Price of Remote Work". NBER Working Paper (29437).

Cahuc, Pierre, Fabien Postel-Vinay & Jean-Marc Robin (2006). "Wage Bargaining with On-the-Job Search: Theory and Evidence". *Econometrica* 74 (2), pp. 323–364.

Chetty, Raj, John N. Friedman, Nathaniel Hendren, Maggie R. Jones & Sonya R. Porter (2025). "The Opportunity Atlas: Mapping the Childhood Roots of Social Mobility". *American Economic Review*, forthcoming.

Chetty, Raj, Michael Stepner, Sarah Abraham, Shelby Lin, Benjamin Scuderi, Nicholas Turner, Augustin Bergeron & David Cutler (2016). "The Association Between Income and Life Expectancy in the United States, 2001-2014". *JAMA* 315 (16), pp. 1750–1766.

Coen-Pirani, Daniele (2010). "Understanding gross worker flows across US states". *Journal of Monetary Economics* 57 (7), pp. 769–784.

Coen-Pirani, Daniele & Holger Sieg (2019). "The impact of the Tax Cut and Jobs Act on the spatial distribution of high productivity households and economic welfare". *Journal of Monetary Economics* 105, pp. 44–71.

Cortés, Patricia, Jessica Pan, Laura Pilossoph, Ernesto Reuben & Basit Zafar (2023). "Gender Differences in Job Search and the Earnings Gap: Evidence from the Field and Lab". *The Quarterly Journal of Economics* 138 (4), pp. 2069–2126.

Costa, Dora L. & Matthew E. Kahn (2000). "Power Couples: Changes in the Locational Choice of the College Educated, 1940–1990". *The Quarterly Journal of Economics* 115 (4), pp. 1287–1315.

DellaVigna, Stefano, Jörg Heining, Johannes F. Schmieder & Simon Trenkle (2021). "Evidence on Job Search Models from a Survey of Unemployed Workers in Germany". *The Quarterly Journal of Economics* 137 (2), pp. 1181–1232.

Dey, Matthew & Christopher Flinn (2008). "Household search and health insurance coverage". *Journal of Econometrics* 145 (1), pp. 43–63.

Din, Alexander & Ron Wilson (2020). "Crosswalking ZIP Codes to Census Geographies: Geoprocessing the U.S. Department of Housing and Urban Development's ZIP Code Crosswalk Files". *Cityscape: A Journal of Policy Development and Research* 22 (1).

Dingel, Jonathan I. & Brent Neiman (2020). "How many jobs can be done at home?" *Journal of Public Economics* 189, p. 104235.

Doepke, Matthias & Michèle Tertilt (2016). "Families in Macroeconomics". In: ed. by John B. Taylor & Harald Uhlig. Vol. 2. Handbook of Macroeconomics. Elsevier, pp. 1789–1891.

Emanuel, Natalia & Emma Harrington (2024). "Working Remotely? Selection, Treatment, and the Market for Remote Work". *American Economic Journal: Applied Economics* 16 (4), pp. 528–59.

Faberman, R Jason, Andreas I Mueller & Ayşegül Şahin (2026). "Job Search, Job Amenities and the Gender Pay Gap". NBER Working Paper (34877).

Findeisen, Sebastian, Sang Yoon (Tim) Lee, Tommaso Porzio & Wolfgang Dauth (2021). "Transforming Institutions: Labor Reallocation and Wage Growth in a Reunified Germany". *Working Paper*.

Flabbi, Luca & James Mabli (2018). "Household Search or Individual Search: Does It Matter?" *Journal of Labor Economics* 36 (1), pp. 1–46.

Galenianos, Manolis & Philipp Kircher (2009). "Directed Search with Multiple Job Applications". *Journal of Economic Theory* 144 (2), pp. 445–471.

Galenianos, Manolis, Philipp Kircher & Gábor Virág (2011). "Market Power and Efficiency in a Search Model". *International Economic Review* 52 (1), pp. 85–103.

Gemici, Ahu (2023). "Family Migration and Labor Market Outcomes".

Gibbs, Michael, Friederike Mengel & Christoph Siemroth (2023). "Work from Home and Productivity: Evidence from Personnel and Analytics Data on Information Technology Professionals". *Journal of Political Economy Microeconomics* 1 (1), pp. 7–41.

Glaeser, Nina, Joan Monras & Milan Quentel (2025). "Spousal Sorting, Spatial Misallocation, and the Gender Wage Gap".

Guler, Bulent, Fatih Guvenen & Giovanni L. Violante (2012). "Joint-Search Theory: New Opportunities and New Frictions". *Journal of Monetary Economics* 59 (4), pp. 352–369.

Guner, Nezih, Remzi Kaygusuz & Gustavo Ventura (2020). "Child-Related Transfers, Household Labour Supply, and Welfare". *The Review of Economic Studies* 87 (5), pp. 2290–2321.

Herkenhoff, Kyle, Gordon Phillips & Ethan Cohen-Cole (2023). "How Credit Constraints Impact Job Finding Rates, Sorting, and Aggregate Output". *The Review of Economic Studies* 91 (5), pp. 2832–2877.

Hosios, Arthur J. (1990). "On the Efficiency of Matching and Related Models of Search and Unemployment". *The Review of Economic Studies* 57 (2), p. 279.

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

Jarosch, Gregor (2023). "Searching for Job Security and the Consequences of Job Loss". *Econometrica* 91 (3), pp. 903–942.

Jayachandran, Seema, Lea Nassal, Matthew J. Notowidigdo, Marie Paul, Heather Sarsons & Elin Sundberg (2024). "Moving to Opportunity, Together". NBER Working Paper (32970).

Jovanovic, Boyan (1979). "Job Matching and the Theory of Turnover". *Journal of Political Economy* 87 (5), pp. 972–990.

— (1984). "Matching, Turnover, and Unemployment". *Journal of Political Economy* 92 (1), pp. 108–122.

Jung, Philip, Philipp Korfmann & Edgar Preugschat (2023). "Optimal regional labor market policies". *European Economic Review* 152, p. 104318.

Jung, Philip & Moritz Kuhn (2019). "Earnings Losses and Labor Mobility over the Life Cycle". *Journal of the European Economic Association* 17 (3), pp. 678–724.

Kaplan, Greg & Sam Schulhofer-Wohl (2017). "Understanding the Long-Run Decline in Interstate Migration". *International Economic Review* 58 (1), pp. 57–94.

Kennan, John & James R. Walker (2011). "The Effect of Expected Income on Individual Migration Decisions". *Econometrica* 79 (1), pp. 211–251.

Kircher, Philipp (2009). "Efficiency of Simultaneous Search". *Journal of Political Economy* 117 (5), pp. 861–913.

Kuhn, Moritz, Iourii Manovskii & Qiu Xincheng (2021). "The Geography of Job Creation and Job Destruction". *NBER Working Paper No. 29399*.

Lange, Fabian & Theodore Papageorgiou (2020). "Beyond Cobb-Douglas: Flexibly Estimating Matching Functions with Unobserved Matching Efficiency". NBER Working Paper (26972).
Lessem, Rebecca (2017). "Mexico–U.S. Immigration: Effects of Wages and Border Enforcement". *The Review of Economic Studies* 85 (4), pp. 2353–2388.
Lise, Jeremy & Jean-Marc Robin (2017). "The Macrodynamics of Sorting between Workers and Firms". *American Economic Review* 107 (4), pp. 1104–1135.
Ljungqvist, Lars & Thomas J. Sargent (1998). "The European unemployment dilemma". *Journal of Political Economy* 106 (3), pp. 514–550.
Lundberg, Shelly & Robert A. Pollak (2003). "Efficiency in Marriage". *Review of Economics of the Household* 1 (3), pp. 153–167.
Maurer, Roy (2017). *Americans Most Often Move for Work*. SHRM. URL: https://www.shrm.org/topics-tools/news/talent-acquisition/americans-often-move-work (visited on 11/21/2025).
Menzio, Guido & Shouyong Shi (2010). "Block Recursive Equilibria for Stochastic Models of Search on the Job". *Journal of Economic Theory* 145 (4), pp. 1453–1494.
— (2011). "Efficient Search on the Job and the Business Cycle". *Journal of Political Economy* 119 (3), pp. 468–510.
Menzio, Guido, Irina A. Telyukova & Ludo Visschers (2016). "Directed Search over the Life Cycle". *Review of Economic Dynamics* 19, pp. 38–62.
Mincer, Jacob (1978). "Family Migration Decisions". *Journal of Political Economy* 86 (5), pp. 749–773.
Pilossoph, Laura & Shu Lin Wee (2021). "Household Search and the Marital Wage Premium". *American Economic Journal: Macroeconomics* 13 (4), pp. 55–109.
Pissarides, Christopher A. (1992). "Loss of skill during unemployment and the persistence of employment shocks". *Quarterly Journal of Economics* 107 (4), pp. 1371–1391.
Piyapromdee, Suphanit (2020). "The Impact of Immigration on Wages, Internal Migration, and Welfare". *The Review of Economic Studies* 88 (1), pp. 406–453.
Postel–Vinay, Fabien & Jean–Marc Robin (2002). "Equilibrium wage dispersion with worker and employer heterogeneity". *Econometrica* 70 (6), pp. 2295–2350.
Rabinovich, Stanislav & Ronald Wolthoff (2022). "Misallocation inefficiency in partially directed search". *Journal of Economic Theory* 206, p. 105559.
Román, Olatz (2025). "The Geography of Jobs and Couple Migration".
Rueda, Valeria & Guillaume Wilemme (2025). "Career Paths with a Two-Body Problem: Colocation and Gendered Professional Crossroads".
Şahin, Ayşegül, Joseph Song, Giorgio Topa & Giovanni L. Violante (2014). "Mismatch Unemployment". *American Economic Review* 104 (11), pp. 3529–3564.
Schaal, Edouard (2017). "Uncertainty and Unemployment". *Econometrica* 85 (6), pp. 1675–1721.
Shimer, Robert (2005). "The Cyclical Behavior of Equilibrium Unemployment and Vacancies". *American Economic Review* 95 (1), pp. 25–49.
Snyder, J. P. (1982). *Map Projections Used by the U.S. Geological Survey*. Geological Survey United States: Geological Survey bulletin. U.S. Government Printing Office.
Venator, Joanna (2025). "Dual-Earner Migration Decisions, Earnings, and Unemployment Insurance".
Wright, Randall, Philipp Kircher, Benoît Julien & Veronica Guerrieri (2021). "Directed Search: A Guided Tour". *Journal of Economic Literature* 59 (1), pp. 90–148.
Zerecero, Miguel (2021). "The birthplace premium".

A Mathematical Appendix

A.1 Kolmogorov Forward Equation

Collect all job-finding rates $\{f_{i,q}(e,s,r)\}$ along with all endogenous separations into μ, let n denote the replacement process of retiring households by new ones, and define $\sigma \equiv \mu + \pi + n$. Then the cross-sectional distribution, $g_t(\boldsymbol{e},\boldsymbol{s},r)$, evolves according to the following differential equation:

$$\frac{\mathrm{d}g_t}{\mathrm{d}t}(\boldsymbol{e}',\boldsymbol{s}',r') = \sum_{\boldsymbol{e},\boldsymbol{s},r} \sigma(\boldsymbol{e}',\boldsymbol{s}',r'|\boldsymbol{e},\boldsymbol{s},r)\, g_t(\boldsymbol{e},\boldsymbol{s},r). \tag{A.1}$$

From (A.1), we may obtain the steady state distribution by setting $\mathrm{d}g = 0$ subject to $\sum g = 1$.

In practice, when estimating the model, we embed the steady-state condition as a constraint into our moments-matching algorithm when computing the distribution of location preference shocks with the smallest total prevalence subject to matching the cross-sectional distribution of households over commuting zones.

A.2 Proof of Lemma 1

We prove a slightly stronger statement. Fix $\boldsymbol{s}$ and let $f : \mathbb{R}^2 \to \mathbb{R}$ be weakly increasing in both arguments. Define

$$\mathcal{F}(\boldsymbol{s}) = \max_{q\in\mathcal{R}} f\left(z_1(\boldsymbol{s},q), z_2(\boldsymbol{s},q)\right).$$

Let $z_1^{\max}(\boldsymbol{s}) \equiv \max_{r\in\mathcal{R}} z_1(\boldsymbol{s},r)$ and $z_2^{\max}(\boldsymbol{s}) \equiv \max_{r\in\mathcal{R}} z_2(\boldsymbol{s},r)$. Then, for any permutation $\{z_1(\boldsymbol{s},r)\}$ holding $\{z_2(\boldsymbol{s},r)\}$ fixed,

$$\mathcal{F}(\boldsymbol{s}) \le f(z_1^{\max}(\boldsymbol{s}), z_2^{\max}(\boldsymbol{s})),$$

where the upper bound is attained by the geographic alignment benchmark.

Lemma 1 follows as the special case where $f(z_1,z_2) = z_1 + z_2$.

A.3 Proof of Proposition 1

Fix $\boldsymbol{e},\boldsymbol{s},r$ and a spouse i such that the state entails migration led by i. By equation (6), as $\eta\to\infty$, spouse i's search effort concentrates on destinations that maximize $\Lambda_{i,q}(\boldsymbol{e},\boldsymbol{s},r)$ (and is split across ties),

$$\lim_{\eta\to\infty} \kappa_{i,q}(\boldsymbol{e},\boldsymbol{s},r) = \begin{cases} \bar{\kappa}_i(\boldsymbol{e},\boldsymbol{s})/|K_i(\boldsymbol{e},\boldsymbol{s},r)|, & q \in K_i(\boldsymbol{e},\boldsymbol{s}) \\ 0, & \text{otherwise,} \end{cases}$$

where K_i is the set of maximizers over destinations,

$$K_i(\boldsymbol{e},\boldsymbol{s}) \equiv \arg\max_q \Lambda_{i,q}(\boldsymbol{e},\boldsymbol{s},r).$$

From (5) and (6), $\Lambda_{i,q}$ is increasing in net migration gains

$$\Delta V_{i,q}(\boldsymbol{e},\boldsymbol{s},r) = V(\boldsymbol{e}^{\text{post},i},\boldsymbol{s},q) - \chi(q|\boldsymbol{s},r) - V(\boldsymbol{e},\boldsymbol{s},r).$$

Hence, $K_i = \arg\max_q \Delta V_{i,q}$. Moreover, under the conditions of the proposition, the ranking of $\Delta V_{i,q}$ is independent of the origin location, so

$$K_i(\boldsymbol{e},\boldsymbol{s}) = \arg\max_q V(\boldsymbol{e}^{\text{post},i},\boldsymbol{s},q).$$

Consequently, conditional on migration, the post-migration continuation value is given by (8).

To complete the proof, note that under the proposition's assumptions, the destination q affects $V(\boldsymbol{e}^{\text{post},i},\boldsymbol{s},q)$ only through the local productivities $z_1(\boldsymbol{s},q), z_2(\boldsymbol{s},q)$ and is weakly increasing in both. Hence, there exists a

composite map $f_{i,\boldsymbol{e},\boldsymbol{s}} : \mathbb{R}^2 \to \mathbb{R}$ such that

$$V(\boldsymbol{e}^{\text{post},i}, \boldsymbol{s}, q) = f_{i,\boldsymbol{e},\boldsymbol{s}}\big(z_1(\boldsymbol{s}, q), z_2(\boldsymbol{s}, q)\big),$$

and $f_{i,\boldsymbol{e},\boldsymbol{s}}$ is weakly increasing in both arguments. Applying a generalized version of Lemma 1 (proved in Appendix A.2) to $f_{i,\boldsymbol{e},\boldsymbol{s}}$ completes the proof.

A.4 Proof of Proposition 2

Let $f_{i,q} = \kappa_{i,q}\lambda(\theta_{\psi_{i,q}})$. For a given value function, search policies in the correlated matching benchmark solve

$$\max_{\{f_{i,q},\omega_q,y_{i,q}\}} \left\{ \sum_{i,q} \left(f_{i,q} - \omega_q\right) \Delta V_{i,q}(\boldsymbol{e}, \boldsymbol{s}, r)\big) + \sum_q \omega_q \Delta V_q^{\text{corr}}(\boldsymbol{e}, \boldsymbol{s}, r) - \sum_{i,q} f_{i,q} y_{i,q} \right\}$$

subject to (2), (3), (7) and (9). Observe that the objective is separable across q and is linear in $\{\omega_q\}$. Thus, for any q, ω_q is optimally set to one of the boundaries in (9), with the upper boundary being strictly optimal if and only if

$$\Delta V_q^{\text{corr}}(\boldsymbol{e}, \boldsymbol{s}, r) > \sum_i \Delta V_{i,q}(\boldsymbol{e}, \boldsymbol{s}, r).$$

This proves Proposition 2.

B Data Appendix

B.1 Data Sources

This appendix describes the data sources that we use in calibrating our model.

American Community Survey (ACS) We construct our ACS sample using the years 2010–2019, restricting attention to individuals aged 18–64, who live in the 48 contiguous states (i.e., excluding Alaska, Hawaii, and Puerto Rico), who are part of a couple (married or cohabiting), and whose partner was also surveyed. 1990 Census Occupation codes are mapped into the occupation classification by Autor & Dorn (2013) using their crosswalk. We deflate wages to 2010 USD using the CPI. To obtain median wages and occupation-shares by commuting zone (CZ), we adopt crosswalks by Autor & Dorn (2013) to aggregate data from Public Use Micro Areas (PUMAs) to the CZ level.

Current Population Survey (CPS) We construct our CPS sample replicating the sample restrictions used for the ACS. 1990 Census Occupation codes are mapped into the occupation classification by Autor & Dorn (2013) using their crosswalk.

Panel Study of Income Dynamics (PSID) We construct our PSID sample to replicate the ACS sample restrictions as closely as possible. We include all existing PSID waves between 2010 and 2019 (2011, 2013, 2017, and 2019). All other sample restrictions used for the ACS are applied identically.

Opportunity Atlas We use CZ-level data from the Opportunity Atlas (Chetty et al. 2025). The Opportunity Atlas draws on various U.S. data sources and aggregates them at the CZ level. For a detailed description of several variables that we draw from the Opportunity Atlas see also Chetty et al. (2016).

Web-scraped data We use web scraping to obtain CZ-level information on local weather and climate conditions, crime rates, walkability scores, measures of beach access and quality, school and hospital quality, and local government expenditures. Specifically, we scrape information published on bestplaces.net, usnews.com, walkscore.com, and watersgeo.epa.gov. The raw data are aggregated at the county and ZIP code levels,

respectively. To aggregate these data to CZs, we use crosswalks by Autor & Dorn (2013) and Din & Wilson (2020). For a description of all web-scraped variables and their sources see Table A.I in Appendix B.3.

Geolocation data To locate CZs in space, we use county centroid geographic coordinate system (GCS) coordinates provided by simplemaps.com. We apply the Albers equal-area projection to transform GCS coordinates into Albers equal-area (AEA) coordinates.[A1] The advantage of using AEA coordinates is that they are Cartesian; i.e., geographic distances can be expressed as Euclidean "straight-line" distances, which allows us to adopt the crosswalk by Autor & Dorn (2013) to aggregate county coordinates to population weighted CZ centroids and further aggregate commuting zones as described in Appendix B.4.

B.2 Measuring Labor Market and Migration Flows

Labor Market Flows We measure labor market flows by using the rotating panel structure of the CPS. The CPS is a monthly survey conducted among a nationally representative sample of households. Leveraging matched individual records for consecutive months, we compute monthly transition frequencies between employment states, conditioning on sex and household occupations.

Migration Flows We measure migration flows using ACS data on where individuals reside when surveyed and their residence one year prior. In particular, the ACS records the Public Use Micro Area (PUMA) individuals reside in when surveyed and the Migration Public Use Micro Area (MIGPUMA) one year prior. MIGPUMAs are constructed from one or multiple PUMAs. We map MIGPUMAs into PUMAs using a crosswalk published by usa.ipums.org, and map PUMAs into CZs using crosswalks developed by Autor & Dorn (2013). In computing across-CZ migration rates, we adopt a conservative approach, counting as cross-CZ migration only a change in residence from a MIGPUMA to a PUMA that have zero overlap in the CZs they intersect with. To obtain spatial distances between migration origin and destination, we use spatial coordinates of population weighted CZ centroids. For differences in population size we use population counts as recorded in the ACS. We aggregate to the PUMA and MIGPUMA level, respectively, using the crosswalks mentioned above together with a crosswalk from counties to CZs also by Autor & Dorn (2013).

B.3 Measuring Amenities

This appendix describes the data inputs and results of regression 12, which we use to convert our CZ-level amenity data into income-equivalent units. Table A.I describes the definition and data source of each amenity in our data. Table A.II summarizes the estimation results of regression (12). Most coefficient estimates are statistically significant at the 1% level. The adjusted R^2 of the regression is 86%.

B.4 Merging Commuting Zones

This appendix describes the algorithm by which we merge CZs that are close geographically and in terms of observed characteristics. Geographic distances between CZs are computed based on AEA coordinates of CZ centroids. We define closeness in terms of all observed CZ characteristics that feed into our calibration; i.e., rents, amenities with and without children (converted to income equivalent units as described in Section 4.2), the local population counts of occupation vector $\boldsymbol{o}$, and gender $\times$ occupation-specific median wages. As a joint measure of geographic proximity and similarity in observables of CZs, we use the weighted Euclidean norm

$$d_{CZ}(r,r') = \sqrt{\frac{1}{J}\sum_{j}\left(\frac{x_j(r)-x_j(r')}{\omega_j}\right)^2},$$

where $x_1(r)$ and $x_2(r)$ are AEA latitude and longitude and $\{x_j(r)\}_{j=3}^{J}$ are the remainder characteristics of CZ r. The weights $\{\omega_j\}$ correspond to the standard deviations of the observed characteristics across CZs, and across all pairwise combinations of CZs for the AEA coordinates.

[A1] We follow Snyder (1982) in using the AEA standard parallels 29.5° and 45.5° north.

Table A.I: Description of amenities data

Variable	Description	Data Source
Population density	Number of people per square mile.	Opportunity Atlas
Local government expenditures	Total local government expenditures per capita, USD per annum.	Opportunity Atlas
Crime rate	Annual crime rate per million people.	Opportunity Atlas
Hospital quality	State's ranking in health care quality. State data is applied to each CZ in the state. Scale: 1–10, 10 being highest quality.	usnews.com
Walkability score	Score based on availability of infrastructure and number of restaurants, bars and coffee shops within 5 minutes walking distance. More choices within a radius yields a higher score. Scale: 0–100, 100 being the highest walkability. Aggregated to CZ level from Zip Code Level.	walkscore.com
No. of top tier beaches	Beaches in the most popular tier, sampled one month before swim season.	watersgeo.epa.gov
Annual precipitation (inches)	Annual inches of precipitation. Aggregated to CZ level from county level.	bestplaces.net
Annual snowfall (inches)	Annual inches of snowfall. Aggregated to CZ level from county level.	bestplaces.net
Average July high temperature (°F)	Aggregated to CZ level from county level.	bestplaces.net
Average school score	Index that combines student proficiency on state standardized tests, SAT/ACT performance, and high school graduation rates. Scale: 1–10, 10 being the best possible score.	greatschools.org
High school dropout rate (income adjusted)	Residual from a regression of high school dropout rates on household income per capita in 2000.	Opportunity Atlas
School expenditure per student	Average expenditures per student in public schools, USD per annum.	Opportunity Atlas

Notes.—We winsorize annual snowfall and local government expenditures at the 99th percentile to reduce the impact of outliers. Missing observations are filled in as follows: For two CZs in New Jersey, school expenditures are filled in with state-level data from nces.ed.gov. For 16 CZs in Florida, and 18 CZs in Illinois, crime rates are filled in with state-level data from ucr.fbi.gov. For one CZ in Connecticut and three CZs in Vermont, high school dropout rates are imputed using linear predictions from a regression of dropout rates on all other amenity variables.

Table A.II: Empirical relationship between amenities and rent

Dependent variable: Annual rent in $	Coefficient estimate	Standard error	P-value
Median household wage (passthrough rate, ϵ)	.161	.008	.000
Population density (people/mile^2)	.501	.082	.000
Local government expenditures ($ per capita)	1.419	.107	.000
Crime rate (per mio. people)	−42.558	23.728	.073
Hospital quality	27.475	2.336	.000
Walkability score	47.912	4.518	.000
No. of top tier beaches	54.192	15.877	.001
Annual precipitation (inches)	−15.408	5.465	.005
Annual snowfall (inches)	−42.696	3.262	.000
Avg. July high temperature (°F)	112.652	20.411	.000
Avg. July high temperature squared	−.065	.012	.000
Share of HHs with children, $\overline{k}_r$ (%)	−16.898	28.544	.554
$\overline{k}_r$ × Avg. school score	4.552	2.591	.079
$\overline{k}_r$ × High school dropout rate (income adjusted)	−1.668	.850	.050
$\overline{k}_r$ × School expenditure per student ($ per annum)	−.001	.001	.405

Notes.—Displayed are coefficient estimates of the relationship between amenities and rents on the CZ level (equation (12)). Hospital quality is recorded on a 1–10 scale, 10 denoting the best possible score. The Walkability score is recorded on a 0–100 scale, 100 denoting highest walkability. The school score is recorded on a 1–10 scale, 10 denoting the best possible score. The regression includes 720 commuting zones (out of a total of 741) for which all amenities and annual rents are observed in our data. The adjusted R^2 of the regression is 0.86.

Equipped with d_{CZ} and some cutoff $\bar{d}$, our algorithm proceeds as follows:

1. Start from the full list of CZs. Sort it in ascending order in terms of population size. Denote the resulting sorted list by $I = (r_0, r_1, ..., r_K)$. Start from $k = 0$.
2. Find $l^* = \underset{l>k}{\operatorname{argmin}}\, d_{CZ}(r_k, r_l)$.
3. If $d_{CZ}(r_k, r_{l*}) \leq \bar{d}$:
   (a) Merge r_{l^*} with r_k (removing r_{l^*} from I). Update the characteristics of the newly merged CZ, $\{x_j(r_k)\}_{j=1}^J$, by summing population sizes and type-$\boldsymbol{o}$ household population counts of r_k and r_{l^*}, and taking population-weighted averages of all other characteristics.
   (b) Iterate through $\{r_m : r_m \in I, m < k\}$ and merge each element r_m satisfying $d_{CZ}(r_m, r_k) \leq \bar{d}$ with r_k, following the steps outlined above to update the characteristics, $\{x_j(r_k)\}_{j=1}^J$, and removing r_m from I.
4. Increase k by 1 and proceed to the next element in I. Repeat steps 2–4 until $k = K$. The resulting I contains the list of merged CZs.

We set the cutoff to $\bar{d} = \frac{1}{3}$, implying that if two commuting zones are identical in all but one characteristic, then the most they can differ in that one characteristic is $\frac{1}{3}$ standard deviation. Starting from 720 CZs with non-missing data, our algorithm delivers a set of 546 merged CZs. The average geographic distance between the centroids of CZs merged by our algorithm is 108 kilometers. The average distance in terms of d_{CZ} is 0.21.

B.5 Additional Figures

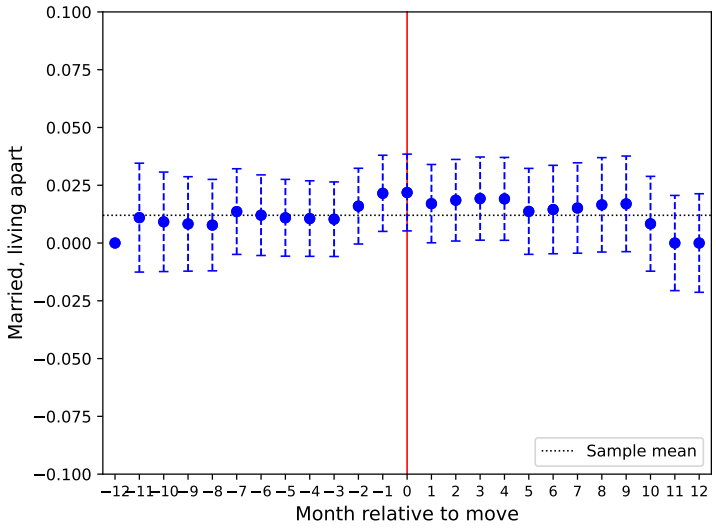

Figure A.I: The plot shows the share of couples who are married but living apart by month relative to migration, conditional on living together t=-12 months prior to migration. The sample is restricted to couples who were living together 12 months prior to migration. The dashed vertical lines indicate 95% confidence intervals. The dotted horizontal line shows the unconditional sample mean. Computed based on married couples in our PSID sample.