

Why Remote Work Stuck: A Structural Decomposition of the Post-Pandemic Equilibrium*

Mitchell Valdes-Bobes

Anna Lukianova

November 3, 2025

[For the latest version, click here.](#)

Abstract

We find that preference shifts, not technology, explain the post-pandemic persistence of remote work. We use a structural equilibrium search model, estimated on 2019 and 2024 microdata, to decompose the 9.1 percentage point rise in remote work. Our decomposition reveals that a "Great Re-valuation" of flexibility by workers accounts for 60% of the increase. Technological improvements (e.g., Zoom, Slack) explain 30%, while changes in skill composition explain 10%. We identify these channels using a novel continuous teleworkability index, which moves beyond binary classifications to capture hybrid arrangements. Our model shows that baseline remote productivity rose 46%, but workers' valuation of remote work increased significantly. Counterfactuals confirm preferences are the primary driver: absent this preference shift, remote work would have risen by only 3.8 percentage points, less than half of the observed increase.

*Valdes-Bobes: University of Wisconsin–Madison. Email: valdsbobes@wisc.edu. Lukianova: University of Wisconsin–Madison. Email: lukianova@wisc.edu.

1 Introduction

Why does remote work remain elevated five years after the pandemic [Barrero et al., 2021]? Conventional wisdom points to technological breakthroughs—Zoom, Slack, cloud infrastructure—that permanently closed the productivity gap between home and office [Davis et al., 2024]. We show this misses the bigger picture. Using a structural model estimated on microdata from 2019 and 2024, we find that shifts in worker preferences, not technology improvements, explain most of the persistent increase in remote work. This challenges the technology-focused narrative dominating policy discussions and reshapes how we think about the future of work.

Our research question is direct: How much of the increase in remote work came from workers valuing flexibility more versus technology making remote work more productive? We answer this by estimating a structural equilibrium model separately for each period and then constructing counterfactuals that isolate each channel. The decomposition reveals a striking imbalance. Starting from a baseline of 6.8% remote work in 2019, the share rose to 15.9% by 2024. Our estimates show that preference shifts alone would have driven remote work to 12.1%—capturing 58% of the total increase. Technology improvements, by contrast, would have produced only 9.5%, accounting for 30% of the change. The remaining 10% reflects compositional shifts in the workforce. Understanding why requires unpacking what we mean by preferences and technology.

The preference channel captures how workers learned about their own valuations during forced mass adoption [Aksoy et al., 2022]. Before 2020, most workers had never experienced sustained remote work. The pandemic served as a giant natural experiment that revealed true preferences—many discovered they valued flexibility far more than previously believed, while others learned they needed in-person interaction. The technology channel, meanwhile, reflects genuine productivity improvements: better videoconferencing, refined cloud collaboration tools, and organizational learning about managing distributed teams [Davis et al., 2024]. These channels interact through workplace arrangements. Workers and firms don’t just negotiate wages—they jointly choose how much time to spend remote versus in-person, trading off productivity against amenity value [Mas and Pallais, 2017]. This joint determination means both channels affect who works where and under what arrangements, making their separation empirically demanding. Recent work by Bagga et al. [2025] examines how amenity shocks drive the dynamic adjustment path—the sustained labor reallocation, surge in vacancies, and vertical Beveridge curve loop following the pandemic. Our focus differs: we decompose the equilibrium level shift to quantify how much preferences versus technology contributed. Despite different approaches, we reach the same conclusion

about preference dominance.

We combine three data sources to measure both outcomes and underlying characteristics. The Current Population Survey provides labor market outcomes and remote work practices for 2019 and 2024. The Occupational Information Network supplies detailed task content and skill requirements for hundreds of occupations. The Bureau of Labor Statistics Occupational Requirements Survey gives us ground truth on which jobs can feasibly be done remotely. Our analysis focuses on employed wage workers aged 25 to 64. What we lack is equally important: we observe two cross-sectional snapshots, not individual transitions, which means we capture equilibrium shifts but not the adjustment paths that got us there.

Separating preferences from technology requires measuring how suitable different jobs are for remote work. We construct a continuous teleworkability index trained on objective survey data using machine learning. Previous research classified occupations as simply “teleworkable” or “not teleworkable” [Dingel and Neiman, 2020]—a binary split that obscures crucial variation [Langemeier and Tito, 2022]. A software engineer and a customer service representative might both be teleworkable, yet their remote productivity differs fundamentally. The engineer might lose collaborative synergies when remote; the representative might not. Our continuous measure captures these gradations, revealing that high-skilled workers systematically sort into highly-amenable occupations—a pattern that proves essential for identifying how skill and technology complement each other.

Four empirical patterns motivate our structural approach. First, remote work rose sharply from 6.8% to 15.9% between 2019 and 2024—a 134% relative increase. Second, a steep educational gradient persists: college graduates work in far more amenable occupations, yet this gap remained remarkably stable through the pandemic. Third, we document persistent wage premiums even after controlling for education and demographics—a concave premium for occupational amenability and an inverted-U premium for actual remote practice that peaks at hybrid arrangements. Fourth, the structure of these premiums shifted dramatically: pre-pandemic markets priced potential for remote work, while post-pandemic markets price actual practice. These facts pose sharp puzzles for simple models.

Each fact demands a specific modeling ingredient. The persistent wage premium for job amenability, even controlling for observables, reveals sorting on unobserved worker skill—high-ability workers must be matching to high-amenable jobs for reasons standard human capital can’t explain. The prevalence of hybrid work, especially the inverted-U wage profile peaking around three days per week remote, rules out simple corner solutions and demands a production technology with genuine interior optima.

The massive dispersion in arrangements among workers in identical occupations with identical education proves that match-specific heterogeneity matters—some workers will commute two hours to avoid remote work while others in the same job work fully remote. Hybrid arrangements deserve special emphasis: they’re not transitional compromises but optimal steady-state choices that balance collaborative productivity gains against flexibility value. Together, these facts require a model with three-dimensional heterogeneity, non-linear production, and explicit bargaining over both wages and arrangements.

We build a search and matching model where workers and firms meet randomly but match selectively based on mutual gains from trade. Three types of heterogeneity drive the equilibrium. Workers differ in skill—their baseline productivity in any match. Firms differ in technology—how suitable their jobs are for remote work, determined by occupation and organizational infrastructure. Individual matches differ in preferences—idiosyncratic tastes for flexibility that get revealed only after the match forms. The production side features two key mechanisms. Remote and in-person work can substitute, but with diminishing returns, which naturally generates hybrid optima. Worker skill and job technology complement each other, so high-skill workers are substantially more productive in high-amenability jobs—this drives positive assortative matching. On the bargaining side, matched pairs jointly choose both wage and work arrangement to maximize their total surplus [Mas and Pallais, 2017], with the wage serving to split the gains and the arrangement balancing productivity against amenity value.

We estimate all model parameters separately for 2019 and 2024, which lets us observe which deep structural features actually changed. The estimation uses moment-matching methods targeting key empirical patterns. Preference parameters—the average taste for flexibility, the dispersion in tastes, and how amenity value curves with remote intensity—come from the observed distribution of work arrangements and how wages vary with actual remote practice. Technology parameters—baseline remote productivity, how worker skill and job technology interact, and how substitutable remote and in-person work are—come from sorting patterns across occupations and wage premiums for job amenability. Estimating both periods separately is crucial: we can’t just assume technology improved and back out preferences, or vice versa. By letting the data speak about both channels simultaneously, we avoid prejudging the answer.

The counterfactual decomposition reveals the dominance of preferences. If only preferences had changed while holding technology fixed at 2019 levels, remote work would have reached 12.1%—a 5.3 percentage point increase that captures 58% of the actual 9.1 point rise. If only technology had improved while holding preferences constant,

remote work would have climbed to just 9.5%—a 2.7 point gain representing 30% of the total shift. Compositional changes in workforce demographics account for the remaining 10%. Technology still matters substantially—we estimate a 46% increase in baseline remote productivity, validating stories about learning-by-doing and network effects. But preferences changed even more.

Preference dominance has important implications. Because the structural shift occurred primarily in workers’ valuations rather than productive capabilities, technology alone cannot account for the persistence we observe. Our findings explain a broader labor market transformation we call the Great Re-Valuation [Aksoy et al., 2022]: markets fundamentally shifted from pricing the scarce option of remote work before the pandemic to pricing actual practice afterward. The flattening premium for potential and steepening premium for practice reflect this maturation—what was once an elite perk became a standard negotiated amenity. This amenity shift has contributed to broader wage dynamics in the post-pandemic labor market [Barrero et al., 2024].

The remainder of the paper proceeds as follows. Section 2 reviews related research on remote work trends, compensating differentials, and equilibrium sorting. Section 3 constructs the teleworkability index and documents key facts. Section 4 presents the structural model. Section 6 describes estimation and identification. Section 6 reports parameters, model fit, and the main decomposition. Section 7 concludes with implications and limitations.

2 Related Literature

Our work speaks to three interconnected literatures: the economics of remote work persistence [see Lee, 2023, for a recent review], the theory and empirics of compensating wage differentials, and structural labor market models with heterogeneous agents. Our analysis provides a quantitative, equilibrium-based decomposition of the forces driving the persistent rise in remote work. We find that preference shifts account for 60% of the increase, technology improvements 30%, and composition changes 10%. This preference-dominant finding contrasts with alternative explanations emphasizing institutional constraints [Adrjan et al., 2023] or compositional sorting forces [Braesemann et al., 2022].

The intellectual foundation for understanding post-pandemic persistence is Barrero et al. [2021], who document that remote work will remain elevated and identify five mechanisms: better-than-expected experiences, investments in infrastructure, diminished stigma, contagion concerns, and technological innovations. We formalize these mechanisms through preference and technology channels in a structural model, moving

beyond survey-based stated preferences to provide *revealed preference* evidence from observed wages, remote work shares, and sorting patterns. This quantifies that preference shifts dominate, providing market-based validation for the [Barrero et al. \[2021\]](#) emphasis on updated worker attitudes. An alternative technology-centric view by [Davis et al. \[2024\]](#) argues that mass adoption unlocked latent productivity through network effects. Our estimation finds a remarkably similar 46% increase in baseline remote productivity, confirming this technology channel is significant. However, by simultaneously estimating both channels, we show that while technological improvements were necessary, preference shifts were dominant—absent technology changes, preferences alone would have driven nearly two-thirds of the observed increase.

Survey evidence from [Aksoy et al. \[2022\]](#) documents that stated preferences for remote work increased during the pandemic, particularly for those with direct exposure. Our structural estimates provide the revealed-preference foundation for this “Great Re-Valuation,” showing a 12% increase in the average taste for flexibility, a 30% increase in the dispersion of tastes across workers, and a 35% decrease in the curvature of amenity value with respect to remote intensity. These changes explain not just higher average remote work but also increased heterogeneity and hybrid arrangements. [Althoff et al. \[2022\]](#) offer an alternative equilibrium selection framework where the pandemic pushed cities between equilibria with stable parameters. Our approach differs fundamentally: we estimate that structural parameters themselves shifted, better explaining within-occupation heterogeneity and systematic wage changes.

Concurrent work by [Bagga et al. \[2025\]](#) explores post-pandemic labor markets through job amenities, developing a frictional search model to explain the adjustment path following the pandemic. They find that an aggregate amenity shock—workers placing higher value on work-from-home arrangements—drives sustained labor reallocation, explaining the rise in job-to-job transitions, the surge in vacancies, the fall in match efficiency, and the Beveridge curve’s vertical loop. Their calibrated model successfully accounts for cross-sectoral patterns of vacancies, quit rates, job-filling rates, and wages during the recovery period. Our work is complementary but distinct in focus: while [Bagga et al. \[2025\]](#) study the *dynamic adjustment path* using time-series variation to infer aggregate shocks, we decompose the *equilibrium level shift* between 2019-2024 into preference versus technology channels using cross-sectional microdata. Despite methodological differences, we reach the same qualitative conclusion—preference shifts dominate technological improvements in explaining remote work persistence. Their finding that amenity shocks drive labor market dynamics validates our quantitative result that preferences account for 60% of the increase, while our explicit separation of preference and technology parameters provides the structural foundation for their

aggregate shock interpretation.

Our analysis of remote work as a job amenity builds on the broader literature valuing working conditions [Maestas et al., 2023] and specifically on Mas and Pallais [2017], who estimate workers willing to forgo 8% of wages for work-from-home options pre-pandemic. We adopt the same framework but provide post-pandemic general equilibrium valuations. Recent evidence suggests dramatic increases: Cullen et al. [2025] find tech workers willing to forgo 25% of compensation, three times the pre-pandemic benchmark. Other post-pandemic estimates find similar, though varied, valuations: Hsu and Tambe [2021] estimate a 7% differential in startup markets, while Lewandowski et al. [2025] find valuation gaps in Poland as high as 24.6% for fully remote work. At the macro level, Barrero et al. [2024] estimate this amenity shift has contributed to a 2.0 percentage point moderation in aggregate wage growth.

Our structural approach contributes nuance by estimating the full preference distribution—the 30% increase in dispersion of tastes explains the coexistence of diverse arrangements, as more workers intensely prefer remote work while others strongly prefer the office. Our most novel empirical contribution documents a structural shift in wage premiums: pre-pandemic markets priced teleworkability potential with steep premiums reflecting scarcity, while post-pandemic markets price actual practice with an inverted-U profile peaking at around 60% remote work (roughly three days per week). This shift from pricing potential to practice reveals market-wide learning about optimal arrangements. A central identification challenge is separating compensating differentials from worker selection. Emanuel and Harrington [2024] use office closures as a natural experiment, finding two-thirds of the productivity gap is selection. We employ a complementary structural strategy where selection parameters (complementarity between worker skill and job technology) are identified by sorting moments, while preference parameters are identified by dispersion in remote work shares and wage profiles, enabling general equilibrium welfare analysis.

Methodologically, we advance teleworkability measurement by building on Dingel and Neiman [2020]’s binary classification. Our approach is motivated by Langemeier and Tito [2022], who demonstrate that such binary "no physical presence" measures show minimal historical change, while communications-based continuous measures (like the one we develop) rose substantially, highlighting the need to capture gradations in remote work feasibility. Our innovations are a continuous index ranging from zero to one, essential for modeling hybrid work and flexible production functions, and a data-driven machine learning approach trained on objective BLS data achieving superior predictive performance.

Our structural framework builds on the standard Diamond-Mortensen-Pissarides

search and matching framework with wage bargaining, incorporating amenities following the hedonic wage framework of [Rosen \[1986\]](#) with recent treatments in [Morchio and Moser \[2024\]](#) and [Albrecht et al. \[2019\]](#). We introduce multi-dimensional heterogeneity—worker skill, firm technology for remote work, and idiosyncratic taste for flexibility—to simultaneously explain systematic sorting of high-skill workers to high-teleworkability jobs and within-match dispersion in arrangements. We also contribute to labor market sorting literature by identifying and quantifying a new dimension: teleworkability. Our counterfactual shows positive assortative matching on this dimension is quantitatively important, as random matching would reduce aggregate remote work by 17%.

3 Data and Stylized Facts

The rise of remote work is a major structural shift in the modern labor market, but rigorous assessment requires moving beyond anecdote and coarse classifications. This section establishes the empirical foundation and vocabulary for our analysis in two parts. First, we introduce a continuous, data-driven index of an occupation’s intrinsic amenability to remote work or teleworkability, we will denote as ψ . We explain how we built it, check that it works, and use it to map out which jobs in the U.S. can actually be done remotely. Second, using this high-resolution lens, we document a set of stylized facts that reveal three core empirical puzzles: *(i)* a persistent, concave wage premium for remote-work potential, even after controlling for observable worker and occupation characteristics; *(ii)* the widespread prevalence of hybrid work arrangements that defy simple corner-solution predictions; and *(iii)* substantial dispersion in work arrangements among observationally identical workers within the same job. These puzzles motivate the structural sorting model with heterogeneous agents developed in the sections that follow.

3.1 Data Sources and Sample Construction

Our empirical analysis combines several data sources, each playing a specific role in building the teleworkability index and understanding remote work patterns. The Current Population Survey (CPS)—a monthly household survey conducted by the U.S. Census Bureau and Bureau of Labor Statistics—serves as our primary data source for labor market outcomes and remote work patterns. The CPS is representative of the civilian employed U.S. population. We focus on two time periods: the pre-pandemic baseline (2018 - 2019), and post-pandemic equilibrium after the labor market had time to adjust (2023 - 2024) . From CPS we obtain information on employment status, occupation,

industry, and demographics. Most importantly for our analysis, it includes data on how much work people actually do remotely (α).

We restrict our sample to employed workers aged 25-64 working for a wage that report positive wages and usual hours worked per week. We exclude workers below 25 years old to focus on individuals who have completed their education and established careers, and we exclude those over 64 to avoid retirement dynamics. Within this age range, we construct four age groups (25-34, 35-44, 45-54, 55-64) to capture lifecycle variation in skill accumulation and job mobility. Statistics reported in this paper are weighted using CPS person weights to ensure national representativeness.

To construct the teleworkability index ψ , we combine two additional data sources. From the Occupational Requirements Survey (ORS), conducted by the Bureau of Labor Statistics we obtain establishment-level of how many workers can work remotely at the occupation level. We consider this survey as the ground truth for training the predictive model. Second, we use the O*NET database, maintained by the U.S. Department of Labor, which contains over 200 standardized measures of occupational characteristics, including task content, work context, skill requirements, and abilities. By training a machine learning model on the intersection of ORS and O*NET data, we can predict telework potential for all occupations in the O*NET system and assign these scores to CPS workers based on their occupation codes.

3.1.1 Measuring Remote Work

An essential variable for this research is the share of work performed remotely, denoted $\alpha \in [0,1]$. This captures the actual intensity of remote work at the worker level. We measure α as the fraction of hours worked remotely to total hours worked. For the period 2022-2024, we measure α directly from newly introduced CPS telework questions. For the pre-2022 period, when these variables were not yet available, we impute α using a gradient-boosted regression model trained on Survey of Income and Program Participation (SIPP) person-month data, with occupation, industry, demographics, and our teleworkability index ψ as predictive features. The SIPP is uniquely suited for training our imputation model because it contains detailed remote work information for the pre-pandemic period, providing a rich baseline for prediction. However, we do not use SIPP directly for our main analysis because of sample selection concerns: the SIPP's panel structure and lower response rates would introduce bias when estimating aggregate labor market patterns, particularly for high-skilled workers who are overrepresented in remote work arrangements. The imputation model achieves strong validation performance ($R^2 = 0.62$, correlation with post-2022 ground truth = 0.71), with occupation and ψ as

the dominant predictors. Full methodological details, including the specific survey variables used, model hyperparameters, cross-validation results, and robustness checks are provided in Data Appendix [B.3](#).

3.2 A Novel Tool to Analyze Remote Work Potential

Before presenting the motivating puzzles of our paper, it is essential to establish the credibility and utility of our key empirical tool: a continuous teleworkability index, denoted as ψ . In this section we introduce the index, detail its construction and validation, and then use it to map the distribution of remote work in the labor market.

3.2.1 A Data-Driven Measure of Occupational Technology

We frame the ψ index not merely as a variable, but as a new analytical “lens” that provides a high-resolution map of the remote work landscape. This lens offers a more nuanced view than the binary classifications introduced by [Dingel and Neiman \[2020\]](#). While foundational, a simple “teleworkable” or “not teleworkable” flag necessarily masks significant heterogeneity. The deficiencies of a binary framework become apparent when examining three professions: a construction worker, a call-center operator, and a software engineer. This binary model accurately classifies the construction worker as non-teleworkable while categorizing both the operator and the engineer as teleworkable. However, the productivity implications and potential for collaborative synergies differ profoundly between the latter two. For the call-center operator, remote productivity may be similar to in-office work. For the software engineer, the collaborative aspect of the job might be impacted by the lack of face-to-face contact, affecting productivity. A continuous measure is required to capture these distinctions.

3.2.2 Construction

Our methodology builds upon the foundational work of [Dingel and Neiman \[2020\]](#), who pioneered the classification of occupations based on their suitability for remote work. Their approach relies on classifying occupations as either teleworkable or not based on specific O*NET dimensions. Our methodology advances this work in two crucial ways.

First, as motivated above, we construct a continuous index rather than a binary flag. This allows ψ to capture the *degree* to which an occupation is amenable to remote work, providing a richer, more realistic measure of occupational technology.

Second, our index is data-driven and predictive, not based on subjective, a priori rules. We train a two-stage machine learning model on objective survey evidence of

remote work feasibility from the Bureau of Labor Statistics’ Occupational Requirements Survey (ORS). The ORS is uniquely suited to serve as our ground truth because it assesses whether telework is permitted as a regular part of performing a job’s “critical function.” The survey’s definition is nuanced: it codes telework as feasible even if some tasks require in-person contact (e.g., a lawyer who can do most work from home but must occasionally appear in court), but not if the core function requires a physical presence (e.g., a delivery driver). This focus on the intrinsic, regular potential for remote work—rather than ad hoc or emergency arrangements—provides a direct, objective measure of occupational amenability that is perfect for training our model. This data-driven approach allows the model to discover complex, non-linear relationships between a rich set of over 200 occupational characteristics from the O*NET database and an occupation’s intrinsic potential for remote work.

The model architecture is designed to address the zero-inflated nature of remote work potential. The first stage employs a Random Forest classifier to distinguish occupations with any remote work potential ($\psi > 0$) from those with none ($\psi = 0$). For occupations identified as “remote-amenable,” a second-stage Random Forest regressor then estimates the continuous share of remote-workable jobs. By training this model on the intersection of the ORS and O*NET datasets, we can identify the most predictive characteristics of remote work suitability. The trained model is then used to impute a remote work potential score for all occupations in the O*NET database. The final output is the comprehensive index, $\psi \in [0, 1]$, representing the predicted fraction of jobs within an occupation that are intrinsically suitable for remote work.

3.2.3 Validation

To move the index from a “black box” to an intuitive and credible measure, it is vital to demonstrate that it captures a meaningful economic concept. This validation proceeds in three steps: assessing the model’s predictive accuracy, interpreting its key drivers, and examining the real-world occupations it identifies at the extremes.

First, we evaluate the predictive performance of the two-stage model on a held-out test set to ensure its technical robustness. The first-stage classifier exhibits high performance, attaining an F1-score of 0.948. This high score is driven by an excellent recall of 0.994, indicating the model successfully identifies nearly all true teleworkable occupations, and a strong precision of 0.906. For the second-stage regressor, the model explains 56.8% of the variance in remote work potential ($R^2 = 0.568$). More intuitively, the Mean Absolute Error (MAE) is a low 0.069, showing that, on average, our model’s prediction for an occupation’s ψ score deviates from the true value by only 6.9 percentage points. Taken

together, these strong performance metrics confirm the model’s robust predictive power and its suitability for creating a credible teleworkability index.

Second, to ensure the model is not just accurate but also intuitive, we examine the O*NET features that drive its predictions. We analyze the MDI (Mean Decrease in Impurity) importance for the top features in both the classifier and regressor stages. This metric quantifies how much each feature contributes to the model’s decision-making accuracy by measuring how effectively it helps to create pure nodes in the decision trees.

The features identified by the model align strongly with economic intuition. For the classifier stage, the ability to work remotely at all is primarily determined by the absence of physical tasks as requirements to perform an occupation. The most important features are all physical requirements that negatively impact teleworkability, such as Stamina, Extent Flexibility, and various forms of physical strength. Conversely, an analytical skill like Written Expression is a positive predictor.


For the regressor stage, among the set of remote amenable jobs, the degree of remote potential is predicted by a similar trade-off. The model assigns a lower ψ score to occupations that still require tasks like Handling and Moving Objects or Assisting and Caring for Others. Conversely, it assigns a higher score to jobs that heavily involve tasks like Working with Computers and Systems Analysis. This confirms that our index is not a statistical black box but is capturing a meaningful economic concept of what makes a job technologically suited for remote work.

Finally, a review of occupations with the highest and lowest ψ scores serves as a powerful real-world validity check. Our index correctly identifies knowledge-intensive roles like lawyers, marketing managers, and financial specialists at the top of the distribution. Conversely, manual or service-based roles requiring physical presence, such as construction workers, assemblers, and food service workers, are correctly identified as having zero potential for remote work.

3.3 The Uneven Terrain of the Modern Labor Market

Next we use our index ψ to paint a broad-strokes picture of the labor market. The central themes that emerge are the scarcity of high-amenability jobs and the powerful concentration of these jobs among specific workers and industries. This descriptive groundwork is not merely background information; it establishes the necessary preconditions for our subsequent puzzles to be economically meaningful. By first establishing that high remote-work potential is scarce and unequally distributed, we highlight the importance of a model for understanding how it is allocated and priced in equilibrium.

Figure 1: Proportion of Workers by Remote Work Amenability



Note: This figure shows the proportion of workers employed in occupations with different levels of teleworkability from 2019 to 2024. The extensive margin reveals that approximately 45% of the workforce is consistently employed in occupations with no intrinsic potential for remote work ($\psi = 0$), highlighting a persistent "remote work desert" in the U.S. labor market.


3.3.1 The Extensive and Intensive Margins

Evidence reveals a "remote work desert." As shown in Figure 1, a large and stable fraction of the U.S. workforce is employed in occupations with no intrinsic potential for remote work ($\psi = 0$). This figure, showing the extensive margin, highlights that the portion of employment in non-teleworkable occupations has been constant around 45% from 2019 to 2024.

For readers seeking more detail on the educational stratification of remote work, a three-panel figure in the appendix presents (i) remote work adoption (α), (ii) teleworkability potential (ψ), and (iii) the implementation gap ($\psi - \alpha$) by education level; see Appendix Figure 11.

Furthermore, for the subset of jobs with any remote potential ($\psi > 0$), the intensive margin reveals that high-amenability is rare. Figure 2 plots the employment-weighted distribution of the ψ index, which is heavily skewed toward lower values. This reinforces the core idea that high-amenability jobs (e.g., $\psi > 0.5$) are the exception, not the rule. This scarcity is consistent feature of the labor market in pre and post pandemic periods, and is a key driver of the sorting patterns we document next.

Figure 2: Distribution of Remote Work Potential for Amenable Jobs



Note: This figure shows the employment-weighted distribution of the teleworkability index (ψ) for jobs with any remote work potential ($\psi > 0$). The distribution is heavily skewed toward lower values, demonstrating that high-amenability jobs (e.g., $\psi > 0.5$) are rare. This scarcity persists across both pre-pandemic and post-pandemic periods.


3.3.2 The Sorting Landscape

Finally, descriptive evidence illustrates the powerful sorting of workers into these scarce, high- ψ jobs. We observe a dramatic educational gradient: while 71.2% of workers with a high school education or less are concentrated in the lowest quartile of remote work potential, 40.7% of college-educated workers are in the highest quartile. This is complemented by strong industry concentration, with Finance and Professional Services overwhelmingly represented in high- ψ occupations.

The sorting of workers into high-teleworkability jobs is not merely a compositional effect but reflects a systematic matching process. Figure 3 presents direct evidence of this relationship through a binned scatter plot of actual remote work share (α) against occupational teleworkability (ψ). The tight, upward-sloping relationship—with a correlation of 0.68—demonstrates that workers in high- ψ occupations are significantly more likely to work remotely. Importantly, the relationship strengthened from 2019 to 2024, with the slope increasing from 0.38 to 0.73, indicating that the pandemic accelerated the alignment between technological potential and actual practice.

This clear occupational divide drives the powerful educational sorting we document,

Figure 3: The Sorting of Remote Work: Practice Follows Potential




Note: Binned scatter plot showing the relationship between actual remote work share (α) and occupational teleworkability (ψ) across all years in the sample. Each point represents an occupation, with bubble size proportional to employment. The tight, upward-sloping relationship (correlation = 0.68) demonstrates strong occupational sorting: high-teleworkability occupations (e.g., Software Developers, Lawyers) exhibit high remote work shares, while low-teleworkability occupations (e.g., Retail Sales Workers, Transportation Managers) remain predominantly in-person. Selected occupations are labeled to illustrate the range.

as illustrated in Figure 4. By plotting the average teleworkability index (ψ) by educational attainment from 2019 to 2024 (Panel A) we illustrate two crucial facts. First, a clear and persistent educational hierarchy exists: post-graduates consistently work in occupations with the highest average ψ , followed by those with bachelor's degrees, some college, and high school degrees. Second, these average levels are remarkably stable, indicating that the pandemic did not fundamentally alter which educational groups have access to remote-amenable occupations, but rather accelerated the use of that pre-existing potential. This descriptive fact of strong assortative matching reveals a striking contrast when compared to actual remote work adoption patterns shown in Panel B. While the technological capacity for remote work remained largely unchanged through the pandemic, the actual utilization of this capacity saw substantial increases, particularly among more educated workers. This divergence between stable potential and changing practice suggests that the rise in remote work cannot be explained by technological changes alone, serving as the natural jumping-off point for our first motivating puzzle.

For readers seeking more detail, a three-panel appendix figure presents (i) remote

Figure 4: Comparing Potential vs. Actual Remote Work by Education Level



Note: Panel A shows the remarkably stable distribution of teleworkability potential (ψ) across education groups from 2019 to 2024. In contrast, Panel B reveals the dramatic increase in actual remote work adoption (α) over the same period, particularly among more educated workers. This divergence between stable technological potential and changing practice provides strong evidence that the rise in remote work reflects more than just technological adaptation.

work adoption (α), (ii) teleworkability potential (ψ), and (iii) the implementation gap ($\psi - \alpha$) by education level; see Appendix Figure 11.

3.4 Three Empirical Puzzles

The sorting patterns documented above raise questions that simple models cannot easily answer. We organize these questions into two core empirical puzzles. First, why do we observe persistent wage premiums for occupational teleworkability even after controlling for education and other observable worker characteristics? This pattern suggests that unobserved skills matter for sorting into remote-amenable jobs. Second, how did the structure of these wage premiums change between 2019 and 2024? We show that the labor market underwent a fundamental shift in how it prices remote work—what we term a "Great Re-Valuation." Together, these patterns motivate the heterogeneous-agent sorting model we develop in Section 4.

3.4.1 Why is Remote-Work Potential Priced in the Market?

The sorting of highly educated workers into remote-amenable jobs is a stark descriptive fact. A simple human capital model would predict that once we control for education and other observable skills, any wage premium associated with an occupation's technology

should disappear. Our first puzzle, therefore, is not simply that sorting occurs, but that a significant wage premium for an occupation’s remote work potential persists after accounting for these observable characteristics. This finding creates a challenge for simple models and points directly toward the importance of sorting on a latent, unobserved skill dimension, which we label h .

We document this persistent premium through a series of Mincer-style wage regressions. The specification progressively introduces controls to isolate the relationship between our teleworkability index, ψ and log wages, including a quadratic term to capture potential non-linearities. The goal of this exercise is not to claim a causal “return to ψ ,” but to provide evidence consistent with a sorting mechanism where higher-skilled workers—both observably and unobservably—sort into high- ψ jobs.

The most comprehensive specification, corresponding to Column (5) in the table, is given by:


$$\log(w_i) = \beta_1\psi_j + \beta_2\psi_j^2 + \mathbf{X}_i'\Gamma + \gamma_s + \delta_t + \zeta_o + \epsilon_i \quad (1)$$

Where $\log(w_i)$ is the logarithm of real wages for individual i , ψ_j is the teleworkability (or remote amenability) index of their occupation j , $\mathbf{X}_i$ is a vector of demographic and educational controls, and γ_s , δ_t , and ζ_o are full sets of fixed effects for state, year, and occupation, respectively.

Our wage regressions systematically document the wage premium associated with teleworkability by progressively adding controls to this specification. The analysis begins with a large, unconditional correlation. The subsequent addition of demographic controls and education reduces the magnitude of the coefficient on ψ , indicating that much of the premium is explained by higher-skilled workers sorting into high- ψ jobs. However, the premium remains large and statistically significant. Finally, the inclusion of detailed industry and broad **occupation group** fixed effects further attenuates the coefficient, yet it remains positive and significant. This final result is particularly telling, as it reveals that even *within* similar occupational categories, a higher potential for remote work is associated with higher wages, suggesting ψ is a priced job amenity that captures technological differences not fully explained by standard occupational classifications.

While our regressions establish the existence of the premium, Figure 5 visualizes its non-linear and concave nature. The binned scatter plot, which shows the relationship after partialling out a full set of demographic and educational controls, clearly illustrates that the wage premium flattens out at the highest levels of remote-work potential. This concavity, captured by the negative and significant coefficient on the squared term (ψ^2) in our regressions, is an important feature that our structural model is designed to

Figure 5: Non-Linear Wage Premium for Teleworkability




Note: Binned scatter plot of residual log wages against the residual teleworkability index (ψ), based on regressions with full demographic and educational controls. The solid red line (quadratic fit) illustrates the concave relationship: the wage premium associated with an occupation's remote work potential is substantial but flattens out at the highest levels of amenability.

rationalize.

A parallel pattern emerges when examining actual remote work practice rather than potential. Figure 6 presents the relationship between residual log wages and residual remote work share (α) after partialling out the same demographic and educational controls. The figure reveals a striking inverted-U shape: wages rise steeply as workers transition from fully in-person ($\alpha = 0$) to hybrid arrangements, reaching a peak around $\alpha \approx 0.6$ (roughly 3 days per week remote), before declining for fully remote workers ($\alpha = 1$). The quadratic fit achieves an R^2 of 0.809 with a highly significant negative coefficient on α^2 (-0.120), confirming strong concavity. This pattern suggests that hybrid arrangements command the highest wage premium, potentially reflecting an optimal balance between collaborative productivity gains from in-person work and amenity value from flexibility. The decline at the fully remote margin may indicate either productivity costs from complete isolation or selection of workers with strong preferences for remote work who accept lower wages.

This empirical evidence of persistent, concave wage premia for both occupational remote work potential (ψ) and actual practice (α) is the primary motivation for our

Figure 6: Wage Premium for Remote Work Practice



Note: Binned scatter plot of residual log wages against residual remote work share (α), controlling for demographics and education. The inverted-U relationship peaks at $\alpha \approx 0.6$, indicating that hybrid arrangements (3 days/week remote) command the highest wage premium. The quadratic fit ($R^2 = 0.809$) captures significant concavity, with wages declining for fully remote workers.

structural model. The model is explicitly designed to rationalize these facts through two key mechanisms:

- **Positive Assortative Matching:** The existence of a premium after controlling for observables motivates the complementarity between latent worker skill (h) and occupational technology (ψ).
- **Non-Linear Production:** The concavity of the premium motivates the use of a CES production function, which allows for diminishing returns and a flexible, non-linear relationship between in-person and remote work.

3.4.2 The Great Re-Valuation: A Structural Shift in the Price of Remote Work

Our analysis thus far has established the key relationships in the post-pandemic labor market. However, our central premise of a “Great Re-Valuation” requires showing that these relationships have fundamentally changed since before the pandemic. To test this, we now formally compare the wage premium structures of the pre-pandemic (2018-2019) and post-pandemic (2023-2024) periods.

To formally test for a structural change in these relationships, we pool the data from the pre-pandemic (2018-2019) and post-pandemic (2023-2024) periods and estimate the following comprehensive wage regression:

$$\begin{aligned}\log(w_{i,j,s,t}) = & (\beta_1\alpha_i + \beta_2\alpha_i^2) + (\gamma_1\psi_j + \gamma_2\psi_j^2) \\ & + \delta_1(D_{post} \times \alpha_i) + \delta_2(D_{post} \times \alpha_i^2) \\ & + \eta_1(D_{post} \times \psi_j) + \eta_2(D_{post} \times \psi_j^2) \\ & + \mathbf{X}_i'\boldsymbol{\Gamma} + \lambda_t + \mu_s + \zeta_o + \epsilon_{i,j,s,t}\end{aligned}\tag{2}$$

where $\log(w_{i,j,s,t})$ is the log real wage for worker i in occupation j , state s , and year t . The variable D_{post} is an indicator equal to one for the post-pandemic years (2023-2024). The model includes quadratic terms for both remote work practice (α_i) and potential (ψ_j), a full vector of individual controls ($\mathbf{X}_i$), and fixed effects for year (λ_t), state (μ_s), and occupation (ζ_o).

The key parameters of interest are the interaction coefficients ($\delta_1, \delta_2, \eta_1, \eta_2$), which capture the change in the wage premium profiles between the pre- and post-pandemic periods. A statistically significant coefficient on these interaction terms provides a formal test of a structural shift. We show in Figures 7 and 8 the predicted wage premium profiles derived from these coefficients, along with their 95% confidence intervals.

The Flattening Premium for Potential (ψ) The first key finding is a structural flattening of the wage premium for an occupation's intrinsic teleworkability. The interaction term between the post-pandemic indicator and the linear ψ term is negative and highly significant (-0.178), while the interaction with the squared term is positive and significant (0.272).


This flattening suggests a maturation of the market. Pre-pandemic, high- ψ was a scarce feature of elite jobs, commanding a large premium. Post-pandemic, with high potential becoming a standard feature in knowledge-based industries, the marginal premium for having a slightly higher ψ score has diminished. The market has shifted from pricing the scarce *option* of remote work to pricing its actual implementation. Economically, this reflects a convergence in wages across the teleworkability distribution: jobs with moderate potential now command wage premiums closer to those of high-potential jobs, reflecting a decline in the scarcity value of the amenity itself.

The Steepening Premium for Practice (α) The other side of this re-valuation is the dramatic steepening of the premium for the actual remote work share. The interaction terms for α are the opposite of those for ψ : the linear interaction is positive and significant

Table 1: Wage Model Comparison: Pre vs Post Pandemic

	log(Wage)		
	(1)	(2)	(3)
Remote Work Intensity (α)	0.342*** (0.033)	0.573*** (0.020)	0.351*** (0.035)
Squared Remote Work Intensity (α)	-0.170* (0.069)	-0.435*** (0.020)	-0.221** (0.074)
Post-Pandemic Indicator $\times$ Remote Work Intensity (α)			0.218*** (0.039)
Post-Pandemic Indicator $\times$ Squared Remote Work Intensity (α)			-0.210** (0.077)
Teleworkability Index (ψ)	0.481*** (0.029)	0.392*** (0.031)	0.522*** (0.026)
Squared Teleworkability Index (ψ^2)	-0.327*** (0.036)	-0.065 (0.038)	-0.314*** (0.033)
Post-Pandemic Indicator $\times$ Teleworkability Index (ψ)			-0.178*** (0.025)
Post-Pandemic Indicator $\times$ Squared Teleworkability Index (ψ^2)			0.272*** (0.034)
Year Fixed Effects	Yes	Yes	Yes
State Fixed Effects	Yes	Yes	Yes
Occupation Fixed Effects	Yes	Yes	Yes
Observations	181,914	177,794	359,708
R^2	0.437	0.410	0.421
R^2_{Within}	0.152	0.137	0.143
Sample:	Pre-Pandemic	Post-Pandemic	Pooled with Interactions
Years:	2018-2019	2023-2024	2018-2019, 2023-2024
Individual Controls:	Yes	Yes	Yes

Figure 7: Wage Premium Profiles (ψ): Pre- vs. Post-Pandemic




Note: The figure shows that before the pandemic, the wage premium for potential was a steep, concave curve (solid blue line), indicating that the market paid a significant premium for jobs that were teleworkable, but this premium saturated at the highest levels of potential. In the post-pandemic era, this relationship has pivoted. The profile is now flatter and more linear (dashed orange line).

(0.218), and the quadratic interaction is negative and significant (-0.210).

The non-overlapping confidence intervals over the majority of the hybrid work range confirm that this structural change is highly statistically significant. The widening of the confidence interval for the pre-pandemic profile at high levels of α reflects the reality that fully remote work was a rare arrangement before 2020. The much tighter post-pandemic interval highlights the emergence of a thick, well-defined market for these arrangements. This steepening reveals that the market now has a much more defined view on the optimal trade-off between productivity and amenities, with a clear compensating differential emerging for workers who are fully remote.

Interpretation: From Potential to Practice The contrasting shifts in the ψ and α profiles tell a coherent economic story. Prior to the pandemic, telecommuting was predominantly a theoretical concept embedded in job descriptions rather than a practical reality. The labor market exhibited a steep ψ premium, consistent with teleworkable occupations commanding higher wages, even though actual remote work was rare (recall that $\bar{\alpha} = 6.8\%$ in 2019 despite $\bar{\psi} = 17.2\%$).

Figure 8: Wage Premium Profiles (α): Pre- vs. Post-Pandemic



Note: The pre-pandemic wage profile for remote work practice was a modest, gentle curve. In contrast, the post-pandemic profile is significantly steeper and more concave. The wage premium now rises sharply as workers move into hybrid arrangements, peaks around a remote work share of 0.65 (roughly 3-4 days per week), and then declines for fully remote workers.

The pandemic forced a massive experiment, and the results were revelatory. Workers learned their true valuation of remote work by experiencing it at scale. Firms learned which tasks could and couldn't be done remotely. This learning caused the market to shift from pricing potential (ψ) to pricing actual practice (α). The flattening of the ψ premium reflects that *being in a teleworkable job* is no longer a scarce luxury; what matters now is *how much you actually work remotely*. The steepening of the α premium reflects a newly calibrated compensating differential: workers now demand higher wages to give up remote work, and firms are willing to pay less to workers who insist on being fully remote (the decline after $\alpha = 0.65$ suggests some productivity loss at the extreme).

This shift has important implications for our structural model. It validates our decision to include both ψ (occupational technology) and $z \cdot \mathcal{A}(\alpha)$ (realized amenity value) as distinct objects. A model with only ψ would miss the heterogeneity in actual arrangements conditional on potential. A model with only α would miss the underlying technological constraints that prevent some workers from working remotely at all. By including both and allowing their relative importance to vary across time periods, we can capture the pre-to-post transition as a shift in how the market prices these two

dimensions.

4 Model

We consider a labor market populated by a continuum of firms and infinitely lived, risk-neutral workers. We abstract from life-cycle and precautionary savings motives and focus only on the trade-offs present in job search and workplace arrangements. The model is populated by heterogeneous agents: workers differ in their skill level, $h \in \mathcal{H}$, which determines their baseline productivity, while firms differ in their remote work efficiency, $\psi \in \Psi$. This parameter primarily reflects occupation-level characteristics that determine a job's suitability for remote work, but it also encompasses firm-specific factors like technological infrastructure and organizational capacity.

The measure of unemployed workers of skill level h is denoted $u(h)$. We assume no on-the-job search; therefore, the aggregate stock of job seekers is $L = \int_{\mathcal{H}} u(h)dh$. Similarly, the measure of vacancies posted for type- ψ jobs is $v(\psi)$, which aggregates to a total of $V = \int v(\psi)d\psi$ job opportunities. This supply of vacancies is determined by a free-entry condition, where firms can create and post jobs at a cost $\kappa(v)$, doing so until the expected value of filling a position equals the cost of posting. Looking for a job is a time-consuming effort, which we characterize by assuming search and matching frictions. A standard constant-returns-to-scale matching function, $M(L, V)$, governs the meeting process between unemployed workers and vacancies. This process determines the outcomes of the search process: the job-finding rate for workers, $p(\theta)$, and the vacancy-filling rate for firms, $q(\theta)$, both of which depend on the aggregate labor market tightness, $\theta = V/L$.

Worker preferences are defined over wage and remote work bundles $\{(w, \alpha)\}_{\mathbb{R}_+ \times [0,1]}$. We assume utility is quasi-linear and increasing in both wages and the amenity value derived from remote work.

The production technology Y depends on the remote work share α , the worker's skill h , and the firm/occupation's remote conduciveness ψ . To allow for a flexible, non-linear relationship between in-person and remote work, we adopt a CES production function. This nests a linear production function where remote and in-person work are perfect substitutes as a special case and allows us to use the data to determine the optimal degree of substitutability (complementarity). The production function is given by:

$$Y(\alpha \mid \psi, h) = Ah \cdot [(1 - \alpha)^\rho + g(h, \psi) \cdot \alpha^\rho]^{1/\rho} \quad (3)$$

Here, Ah is the Total Factor Productivity (TFP) of the match, which is linear in worker

skill. The parameter ρ governs the elasticity of substitution between in-person work ($1 - \alpha$) and remote work (α), which is given by $\sigma = 1/(1 - \rho)$.

Remote work productivity is scaled by a relative efficiency adjustment factor, $g(h, \psi)$, which we specify with a flexible functional form:

$$g(h, \psi) = \psi_0 \cdot h^\phi \cdot \psi^\nu \quad (4)$$

In this specification, ψ_0 is a baseline technology parameter for remote productivity. We interpret this parameter as the economy-wide technological capability for remote work; think of the adoption of the internet as a permanent increase in ψ_0 in the 1990s and the recent improvement of video-conference software like Zoom or Microsoft Teams as further increases of ψ_0 . The parameters ϕ and ν are the output elasticities with respect to worker skill and firm remote work efficiency, respectively. This functional form implies a **complementarity** between worker skill and firm remote efficiency, which is a central force driving assortative matching in the model.

4.1 Benchmark Model with Homogeneous Preferences

To build intuition for the core trade-offs governing workplace flexibility, we first analyze a simplified benchmark model where the choice of work arrangement is deterministic. This approach allows us to isolate the fundamental economic forces driven by technology and preferences before we introduce additional sources of heterogeneity.

Assume that worker preferences are given by:

$$u(w, \alpha) = w + \bar{z} \cdot \mathcal{A}(\alpha) \quad (5)$$

where $\mathcal{A}(\alpha)$ is the common valuation that workers place on $\alpha \in [0, 1]$ amount of remote work flexibility and $\bar{z}$ is just a scaling factor which we assume is common for all workers. The firm's per-period profit is its output net of the wage paid, $\pi(h, \psi \mid w, \alpha) = Y(h, \psi \mid \alpha) - w$. As the wage, w , is a pure intra-match transfer, it cancels out when we consider the joint surplus of the match. This leaves the surplus dependent only on the match's total output and the worker's non-pecuniary amenity value.

We assume the wage is determined by generalized Nash bargaining, which implies an efficient contract. The remote work share, α , is therefore chosen jointly to maximize the total surplus generated by the match:

$$\pi(h, \psi) = \max_{\alpha \in [0, 1]} \left\{ Y(\alpha \mid \psi, h) + \bar{z} \cdot \mathcal{A}(\alpha) \right\} \quad (6)$$

The solution to this problem, $\alpha^*(\psi, h)$, is found by balancing the marginal productivity of remote work against its marginal amenity value. The first-order condition for an interior solution is:

$$\frac{\partial Y}{\partial \alpha} + \bar{z} \cdot \frac{\partial A}{\partial \alpha} = 0 \quad (7)$$

This condition states that the optimal hybrid arrangement occurs where the marginal gain in production from an additional unit of in-person time is exactly equal to the marginal amenity value lost. The optimal work arrangement partitions the market into three distinct regimes based on the firm's remote efficiency ψ and the worker's skill h :

1. **Full In-Person** ($\alpha^* = 0$): This occurs in matches where the marginal productivity of the first increment of remote work is negative and so large that it outweighs the marginal amenity gain. This is typical for low-skill workers in low-efficiency occupations.
2. **Full Remote** ($\alpha^* = 1$): This occurs in matches where, even at full remote work, the marginal productivity of remote time still outweighs the diminishing marginal amenity value. This is characteristic of high-skill workers in highly efficient remote occupations.
3. **Hybrid Work** ($0 < \alpha^* < 1$): This emerges for the intermediate matches where a clear trade-off exists. The match chooses an interior solution that perfectly balances the marginal trade-off between production and the amenity value.

This deterministic model provides a sharp prediction: for any given worker-firm type (h, ψ) , there is a single, optimal work arrangement. However, this stark segmentation is at odds with the smooth distribution of work arrangements observed in the data, where we see a wide variety of choices even within the same occupation and for workers with similar observable skills. This motivates the introduction of idiosyncratic, heterogeneous preferences, which allows for a richer and more realistic pattern of matching.

4.2 Full Model with Heterogeneous Preferences

The variability in work arrangements is explained by an idiosyncratic component in the worker-firm match. We posit that each potential match has an intrinsic non-pecuniary amenity value associated with flexible work arrangements, which is uncertain ex ante but resolved once the match forms. This heterogeneity is captured by a match-specific preference parameter.

Similar to the deterministic case, a worker's utility in a given match is the sum of their wage and the non-pecuniary amenity value they derive from that arrangement:

$$u(w, \alpha; z) = w + z \cdot \mathcal{A}(\alpha) \quad (8)$$

Here, z is the match-specific preference parameter, which represents the idiosyncratic value of flexibility in this particular worker-firm pairing. A high z signifies a match where the worker strongly benefits from remote work amenities relative to alternatives. The function $\mathcal{A}(\alpha)$ is a deterministic component that governs how the amenity scales with the remote share α .

Analogous to the deterministic model, the total joint surplus for a given arrangement is the sum of the firm's profit and the worker's utility. The wage remains a pure transfer and cancels out, leaving the surplus dependent only on the physical output and the match-specific non-pecuniary utility:

$$\pi(h, \psi; z) = Y(\alpha \mid \psi, h) + z \cdot \mathcal{A}(\alpha) \quad (9)$$

We assume that the match-specific preference parameter z is drawn from a log-normal distribution across potential matches in the population:

$$z \sim \text{LogNormal}(\mu_z, \sigma_z) \quad (10)$$

This flexible specification allows for a realistic representation of match heterogeneity. The parameter μ_z denotes the mean of the normal distribution, while σ_z represents the standard deviation that influences match value dispersion. This distributional assumption is empirically motivated—for instance, workers may apply to two otherwise identical jobs, but one requires a 2-hour commute while the other is a 10-minute bike ride; in the former case, the effective preference for remote work (via z) rises substantially to compensate for the disamenity. It allows us to capture the full range of match-specific preferences observed in the data.

Before the match forms and the specific preference type z is revealed, the firm and worker jointly maximize the expected value of their potential match. In plain terms, when a firm considers hiring a worker (or a worker considers a job), neither knows the realized z for that pairing, but they can evaluate it based on the average across possible realizations. This “inclusive value” accounts for the option to choose the best possible work arrangement after z is revealed upon match formation (e.g., during bargaining). The ex-ante joint maximization problem is:

$$\pi(h, \psi) = \mathbb{E}_z \left[\max_{\alpha \in [0,1]} \{Y(\alpha | \psi, h) + z \cdot \mathcal{A}(\alpha)\} \right] \quad (11)$$

4.2.1 Solving for the Optimal Work Arrangement

For a given realization of z , the optimal work arrangement $\alpha^*(z; h, \psi)$ is determined by the first-order condition that balances the marginal productivity loss from remote work against the marginal amenity gain:

$$\frac{\partial Y(\alpha | h, \psi)}{\partial \alpha} + z \cdot \frac{\partial \mathcal{A}(\alpha)}{\partial \alpha} = 0 \quad (12)$$

We can rewrite this condition by defining a **cutoff function** $Z^*(\alpha; h, \psi)$ that gives the minimum taste shock required for a worker to choose a given remote work share α :

$$Z^*(\alpha; h, \psi) = - \frac{\partial Y(\alpha | h, \psi) / \partial \alpha}{\partial \mathcal{A}(\alpha) / \partial \alpha} \quad (13)$$

This cutoff function maps each remote work share to the preference intensity that would make a worker exactly indifferent at that point. Workers with $z > Z^*(\alpha)$ prefer more remote work than α , while those with $z < Z^*(\alpha)$ prefer less.

The optimal choice falls into one of three regions:

1. **Full In-Person** ($\alpha^* = 0$): For $z < Z^*(0; h, \psi)$, the marginal amenity value is too low to compensate for any productivity loss, leading to full in-person work.
2. **Interior Solution** ($0 < \alpha^* < 1$): For $Z^*(0; h, \psi) \leq z \leq Z^*(1; h, \psi)$, workers choose $\alpha^*(z)$ by inverting the cutoff function to satisfy the first-order condition.
3. **Full Remote** ($\alpha^* = 1$): For $z > Z^*(1; h, \psi)$, the amenity value is so high that workers choose maximum remote work even at the boundary.

4.2.2 Computing the Inclusive Value

The inclusive value is computed by integrating over all possible taste realizations, accounting for the optimal choice in each region:

$$\begin{aligned} \pi(h, \psi) = & \int_0^{Z^*(0)} Y(0 | h, \psi) dF(z) \\ & + \int_{Z^*(0)}^{Z^*(1)} [Y(\alpha^*(z) | h, \psi) + z \cdot \mathcal{A}(\alpha^*(z))] dF(z) \\ & + \int_{Z^*(1)}^{\infty} [Y(1 | h, \psi) + z \cdot \mathcal{A}(1)] dF(z) \end{aligned} \quad (14)$$

where $F(z)$ is the cumulative distribution function of the log-normal distribution. This integral is evaluated numerically using precomputed grids for the functions $Y(\alpha)$, $\mathcal{A}(\alpha)$, and $Z^*(\alpha)$ for each worker-firm type pair (h, ψ) .

We define the flow surplus of the match, $s(h, \psi)$, as this value net of the unemployed worker's outside option, which is their value of home production b :

$$s(h, \psi) = \pi(h, \psi) - b \quad (15)$$

We assume the value of home production is constant across skill levels, representing a baseline level of utility from non-market activities.

The distribution of work arrangements that emerges from this model is determined by the distribution of taste shocks. For a given match type (h, ψ) , the probability density of observing remote work share α is:

$$p(\alpha | h, \psi) = f(Z^*(\alpha; h, \psi)) \cdot \left| \frac{\partial Z^*(\alpha; h, \psi)}{\partial \alpha} \right| \quad (16)$$

where $f(\cdot)$ is the probability density function of the log-normal distribution. This formula accounts for the Jacobian of the transformation from the taste shock distribution to the distribution over work arrangements. The parameter σ_z governs how dispersed choices are around the “typical” arrangement, with larger σ_z leading to more heterogeneity in observed work arrangements for similar worker-firm pairs.

4.3 Equilibrium Characterization

A steady-state equilibrium is characterized by a set of value functions for workers and firms, optimal vacancy posting by firms, and worker flows that are balanced. These components are mutually consistent and determine the aggregate state of the labor market.

4.3.1 Value Functions and Surplus

The equilibrium is defined by the lifetime values for workers and firms in different states. We begin by defining the value of an ongoing match and the value of unemployment, and from these, we derive the Bellman equation for the match surplus, $S(h, \psi)$, which will be shown to be the central object that characterizes the equilibrium.

The joint value of a match, $J(h, \psi)$, is the present discounted value of all future returns. It is composed of the current period's expected flow surplus, $s(h, \psi)$, plus the discounted continuation value. With probability $(1 - \delta)$, the match survives, and with probability δ ,

it is destroyed and the worker's value reverts to that of unemployment, $U(h)$.¹ This gives the Bellman equation:

$$J(h, \psi) = s(h, \psi) + \beta [(1 - \delta)J(h, \psi) + \delta U(h)] \quad (17)$$

The value of unemployment for a worker of type h , $U(h)$, consists of the current flow benefit, b , plus the expected value from job search. With probability $p(\theta)$, the worker contacts a firm, and with probability $(1 - p(\theta))$, they remain unemployed. The value of a new match to a worker, $W(h, \psi')$, is determined by our **Nash bargaining** assumption, which dictates that the worker receives their outside option, $U(h)$, plus a share ξ of the total match surplus, $S(h, \psi')$. A match is only formed if the surplus is positive. This gives the Bellman equation for unemployment:

$$U(h) = b + \beta \left[p(\theta) \mathbb{E}_{\psi'} [\max\{W(h, \psi'), U(h)\}] + (1 - p(\theta))U(h) \right] \quad (18)$$

Substituting $W(h, \psi') = U(h) + \xi S(h, \psi')$ and simplifying yields:

$$U(h) = b + \beta U(h) + \beta p(\theta) \xi \int \max\{0, S(h, \psi')\} d\Gamma_v(\psi') \quad (19)$$

The total match surplus is the net value created by the match, defined as $S(h, \psi) \equiv J(h, \psi) - U(h)$. We can derive its Bellman equation by subtracting the equation for $U(h)$ from the one for $J(h, \psi)$, which after rearranging highlights that the worker's expected gain from search acts as an effective opportunity cost for the match:

$$S(h, \psi) = s(h, \psi) + \beta(1 - \delta)S(h, \psi) - \beta p(\theta) \xi \int \max\{0, S(h, \psi')\} d\Gamma_v(\psi') \quad (20)$$

Solving for $S(h, \psi)$ gives the final expression:

$$S(h, \psi) = \frac{s(h, \psi) - \beta p(\theta) \xi \int \max\{0, S(h, \psi')\} d\Gamma_v(\psi')}{1 - \beta(1 - \delta)} \quad (21)$$

This derivation highlights a key insight: the surplus equation is the only object needed to solve for the equilibrium. Because the term $\max\{0, S(h, \psi')\}$ appears inside the integral, the sign of the surplus itself determines the set of viable matches in equilibrium. A negative surplus means the parties are better off separated, and no match is created. Therefore, $S(h, \psi)$ is a sufficient statistic that fully encodes the equilibrium matching decisions.

¹We assume that free entry makes the value for a firm of a discontinued match equal to zero.

4.3.2 Vacancy Creation and Market Tightness

The number of vacancies is determined by firms' profit maximization. Firms post vacancies for a type- ψ job until the marginal cost of posting, $\kappa'(v)$, equals the expected marginal benefit. The benefit of posting depends on the probability of filling the vacancy, $q(\theta)$, and the firm's expected share of the surplus, $(1 - \xi)S(h, \psi)$, averaged over the distribution of workers it might meet. This gives the vacancy creation condition:

$$c'(v(\psi)) = q(\theta)(1 - \xi) \int_{\mathcal{H}} \max\{0, S(h, \psi)\} \frac{u(h)}{L} dh \quad (22)$$

This set of decisions, aggregated across all firm types, endogenously determines the total stock of vacancies V and thus the equilibrium market tightness θ .

4.3.3 Steady-State Flows

In a steady-state equilibrium, the flows of workers between employment and unemployment are balanced. The total number of workers who lose their jobs must equal the total number of unemployed workers who find new, acceptable jobs. This condition, $\delta \cdot N_{\text{emp}} = p(\theta) \cdot N_{\text{unemp}} \cdot \mathbb{P}(\text{Accept})$, where N denotes the mass of workers, closes the model by determining the equilibrium distributions of employed and unemployed workers for each skill type, $n(h, \psi)$ and $u(h)$.

4.4 Wage Determination

With the expected surplus of the match determined, the wage serves as the transfer that divides the realized proceeds. It is a **contingent contract** where the payment depends on the realized work arrangement, α^* , which is chosen after the worker's idiosyncratic preference, z , is revealed.

To derive the wage, we start with the worker's utility in an ongoing match, $u(w, \alpha^*) = w + z \cdot A(\alpha^*)$. The Bellman equation for the worker's value, $W(h, \psi)$, is:

$$W(h, \psi) = (w + z \cdot A(\alpha^*)) + \beta(1 - \delta)W(h, \psi) + \beta\delta U(h) \quad (23)$$

Rearranging this asset-pricing equation, we can solve for the wage that must be paid to deliver the lifetime value $W(h, \psi)$:

$$w = (1 - \beta(1 - \delta))W(h, \psi) - \beta\delta U(h) - z \cdot A(\alpha^*) \quad (24)$$

From our Nash bargaining assumption, we know the worker's equilibrium lifetime value is $W(h, \psi) = U(h) + \xi S(h, \psi)$. Substituting this into the wage equation gives the

final expression, which separates the wage into two distinct components:

$$w^*(h, \psi; \alpha^*) = \underbrace{((1 - \beta(1 - \delta)) [U(h) + \xi S(h, \psi)] - \beta \delta U(h))}_{\text{Base Wage}} - \underbrace{z \cdot A(\alpha^*)}_{\text{Amenity Value Deduction}} \quad (25)$$

1. **Base Wage:** The monetary payment required to deliver the worker their bargained share of the match's total expected surplus, $\pi(h, \psi)$. This term is determined at the start of the match and is constant with respect to the realized α .
2. **Amenity Value Deduction:** A state-contingent deduction from the base wage. Since the worker receives a non-pecuniary amenity value of $z \cdot A(\alpha^*)$, their monetary wage is lowered by exactly this amount. This ensures the division of the total surplus remains consistent with the parties' bargaining powers in every state of the world. This term is zero for a fully in-person worker ($\alpha = 0$).

In the computational implementation, wages are calculated ex-post after solving for the equilibrium distributions of workers and firms, using the converged values of $S(h, \psi)$ and $U(h)$ from the fixed-point iteration.

This framework thus provides a complete mapping from our structural parameters (θ) to the equilibrium distributions of wages, remote work shares (α), and job types (ψ). The model generates equilibrium outcomes that can be directly compared to their empirical counterparts: observed wage distributions, remote work adoption patterns, and the sorting of workers across occupations. This mapping allows us to take the model to the data using Simulated Method of Moments, which we describe next.

5 Estimation Strategy

To identify the model's 11 structural parameters, we use a Simulated Method of Moments (SMM) strategy. The credibility of our identification rests on mapping specific groups of parameters to distinct, informative moments in the data—a strategy we detail explicitly in Table 2. This section describes our estimation approach, the moment-to-parameter mapping that ensures identification, and our computational implementation.

5.1 Overview of SMM Approach

The model contains 11 structural parameters that govern preferences, technology, and the skill distribution. Let $\theta \in \Theta \subset \mathbb{R}^{11}$ denote the parameter vector. For a given θ , we

solve the model to obtain equilibrium distributions and value functions, then simulate a large sample of workers to compute model-implied moments $m_{\text{model}}(\theta)$. We choose $\hat{\theta}$ to minimize the weighted distance between model and data moments:

$$\hat{\theta} = \arg \min_{\theta \in \Theta} [m_{\text{model}}(\theta) - m_{\text{data}}]' W [m_{\text{model}}(\theta) - m_{\text{data}}] \quad (26)$$

where m_{data} is the vector of 19 empirical moments and W is a positive definite weighting matrix. We use the inverse of the diagonal variance matrix estimated from 500 bootstrap samples, giving more weight to precisely measured moments.

5.1.1 Moment Computation

For each candidate parameter vector θ , we follow a three-stage procedure to compute model-implied moments. First, we solve for the equilibrium using fixed-point iteration to find market tightness θ^* , value functions $U(h)$ and $S(h, \psi)$, and the employment distribution $n(h, \psi)$. The solver uses adaptive step sizes and typically converges in 5-10 seconds with tolerance 10^{-6} .

Second, we simulate worker choices by drawing (h_i, ψ_i) for $N = 100,000$ workers from the equilibrium employment distribution $n(h, \psi) / \sum n$. For each worker, we draw a taste shock $z_i \sim \text{LogNormal}(c_0, \sigma)$ and solve for the optimal remote work share α_i^* using the first-order condition from the worker's optimization problem. We then calculate the wage using equation (25), which accounts for both the bargained share of production and the compensating differential for remote work amenities.

Third, we compute moments using the simulated data $\{(h_i, \psi_i, z_i, \alpha_i^*, w_i)\}_{i=1}^N$ along with demographic characteristics from a pre-loaded scaffold. This stage involves computing distributional moments such as means and variances, running wage and sorting regressions with fixed effects, and calculating subsample statistics for non-teleworkable jobs.

To reduce simulation noise and improve computational efficiency, we fix the random draws (u_h, u_ψ, u_z) across all evaluations using a pre-generated scaffold of 100,000 observations with demographic characteristics sampled from the 2024 CPS. This approach ensures that differences in moment values across parameter vectors reflect genuine economic effects rather than simulation noise.

5.2 Target Moments and Identification

Credible identification requires linking each parameter to moments that provide sharp information about its value. We target 19 empirical moments organized into four groups:

Table 2: Identification: Mapping Moments to Parameters

Parameter(s)	Target Moment	Intuition
<i>Worker Skill Distribution</i>		
a_h, b_h	Mean of log wages	Central tendency
	Variance of log wages	Wage dispersion
<i>Preferences (Taste Shocks)</i>		
μ_z	Wage gap (remote vs. in-person)	Location of taste distribution
χ	Share fully remote ($\alpha = 1$)	Preference curvature
σ_z	Share hybrid ($0 < \alpha < 1$)	Taste dispersion
<i>Production Technology</i>		
ν	Wage- ψ slope	Firm-remote complementarity
ψ_0	Wage premium (high- ψ firms)	Productivity scale
ϕ	$\mathbb{E}[\alpha \mid \psi_{\text{high}}] - \mathbb{E}[\alpha \mid \psi_{\text{low}}]$	Skill-remote sorting
<i>Search Frictions</i>		
κ_0	Market tightness (V/U ratio)	Vacancy cost

Notes: This table summarizes the primary identifying moments for each structural parameter. Each parameter is pinned down by moments that are most sensitive to changes in that parameter, ensuring credible identification. The model is estimated using 19 target moments organized into distributional moments, wage regression coefficients, sorting patterns, and search friction measures.

distributional moments that pin down preference and skill distributions, regression coefficients that identify technology parameters through wage and sorting patterns, subsample moments for non-teleworkable jobs that validate technological constraints, and an external market tightness measure that disciplines search frictions. Table 2 summarizes the mapping between parameters and their primary identifying moments.

5.2.1 Distributional Moments

The first group consists of nine unconditional distributional moments that anchor the model's core distributions. The moments for the wage distribution—specifically the mean log wage, variance of log wages, and the p90/p10 ratio—primarily identify the TFP scale parameter A_1 and the skill distribution parameters (a_h, b_h) . The mean log wage pins down the overall productivity level, while the variance and p90/p10 ratio capture wage dispersion stemming from skill heterogeneity and differences in match quality. For the remote work distribution, the mean remote work share ($\bar{\alpha}$) identifies the central tendency of the taste distribution (parameter c_0), while its variance identifies taste heterogeneity (σ) and amenity curvature (χ). The shares of workers in fully in-person, hybrid, and fully remote arrangements capture the discrete classification of work arrangements, helping to identify the tails of the taste distribution and how amenity curvature (χ) affects concentration at corner solutions. Finally, the mean teleworkability

$(\bar{\psi})$ reflects the job distribution and serves as a consistency check, as ψ is taken directly from the data rather than being estimated.

5.2.2 Wage Regression Moments

We estimate a unified wage regression to capture how wages vary with job characteristics:

$$\log w_i = \beta_1 \psi_i + \beta_2 \psi_i^2 + \beta_3 \alpha_i + \beta_4 \alpha_i^2 + X_i' \gamma + \text{FE}_{\text{state}} + \text{FE}_{\text{industry}} + \epsilon_i \quad (27)$$

where X_i includes age, experience (quadratic), sex, race, and education. This single regression yields four moments. The first two, the coefficients on the flexibility premium (β_1, β_2) , capture how wages vary with job teleworkability. They primarily identify the remote productivity parameters (ψ_0, ν, ϕ) , which determine how productive remote work can be as a function of job characteristics and worker skill; the linear term β_1 reflects the average wage premium for teleworkable jobs, while the quadratic term β_2 captures diminishing returns. The other two moments, the coefficients on remote work wage effects (β_3, β_4) , capture how wages vary with actual remote work intensity. The linear term β_3 primarily reflects the production-technology trade-off between in-person and remote work (governed by ρ), while the quadratic term β_4 identifies the amenity curvature χ through compensating wage differentials.

5.2.3 Sorting Moment

To capture how workers sort into remote work based on job characteristics, we estimate:

$$\alpha_i = \lambda \psi_i + X_i' \delta + \text{FE}_{\text{state}} + \text{FE}_{\text{year}} + u_i \quad (28)$$

The coefficient λ measures the degree of positive assortative matching between teleworkable jobs and remote work adoption. This moment helps identify the interaction between preference and technology parameters, particularly how taste heterogeneity σ and remote productivity ν jointly determine sorting patterns.

5.2.4 Non-Teleworkable Job Moments

The fourth group of moments focuses on jobs with zero teleworkability ($\psi = 0$), which provide sharp identification for the model's technological parameters. We target three moments for this subsample. First, the mean remote work share for these jobs ($\bar{\alpha}|\psi = 0$) should be close to zero, which validates the model's prediction that remote work is infeasible in non-teleworkable jobs. Second, the wage premium for teleworkable jobs over non-teleworkable ones ($\bar{w}|\psi > 0 - \bar{w}|\psi = 0$) identifies the productivity gains associated

with teleworkability. Finally, the share of jobs with zero teleworkability (40.5% in the data) provides crucial variation in job characteristics that helps to identify the production function.

5.2.5 External Moment

Finally, we include one external labor market moment to discipline the search and matching components of the model. We target the aggregate market tightness ($\theta = V/U$), taken from JOLTS data. This moment identifies the vacancy posting cost parameter, κ_0 , through the model's free-entry condition, as shown in Equation (22). We do not separately target the job-filling rate, as it is mechanically linked to market tightness via the matching function.

5.3 Parameter Groups and Bounds

The preference parameters govern worker utility from remote work. The taste distribution for the idiosyncratic preference shock, z , is assumed to be LogNormal and is controlled by a location parameter, c_0 , and a scale parameter, σ . We estimate c_0 directly with bounds of $[-3, 5]$, while σ is estimated using an exponential transformation to ensure positivity. The amenity curvature parameter, χ , governs the concavity of the amenity value in the remote work share and allows for a wide range of shapes, from nearly linear ($\chi \approx 0.08$) to highly concave ($\chi \approx 20$) amenities.

The technology parameters define the production function. The productivity of remote work is determined by the function $g(h, \psi) = \psi_0 \cdot h^\phi \cdot \psi^\nu$, which includes a baseline remote productivity parameter, ψ_0 , and elasticities with respect to worker skill, ϕ , and job teleworkability, ν . All three parameters are constrained to be positive while allowing for a wide range of elasticities. The substitution elasticity, ρ , governs the substitutability between in-person and remote work. We estimate $\rho \in (0, 1)$ using a bounded transformation, allowing the data to determine whether in-person and remote work are complements ($\rho < 0.5$) or substitutes ($\rho > 0.5$). Finally, the overall TFP scale factor, A_1 , is estimated with bounds of $[5, 500]$.

The final set of parameters governs market-level forces and the skill distribution. The vacancy cost parameter, κ_0 , is the scale parameter for the cost of posting vacancies and is estimated with bounds of $[0.1, 5000]$. We assume the skill distribution follows a Beta distribution, characterized by two shape parameters, a_h and b_h . We use transformations to ensure both parameters are positive and to better explore their ratio.

5.4 Computational Implementation

Estimating our structural model presents computational challenges. The objective function is expensive to evaluate (30-60 seconds per call) due to solving equilibrium conditions and simulating 100,000 workers. The parameter space is high-dimensional with 11 parameters, and the objective surface exhibits multiple local minima. We address these challenges through a two-stage hybrid optimization approach that combines global exploration with local refinement. The first stage uses a genetic algorithm distributed across a computing cluster to identify promising regions of the parameter space. The second stage applies a gradient-free trust-region method to achieve high-precision convergence from the best candidate.

5.4.1 Global Search Phase

Parameter estimation proceeds in two stages, beginning with a global search to thoroughly explore the parameter space. We employ a genetic algorithm (GA) that uses tournament selection, blend crossover, and adaptive mutation to explore approximately 500,000 parameter vectors across 1000 generations. The algorithm is distributed across a 48-core computing cluster using MPI (Message Passing Interface), with each worker independently evaluating candidate parameter vectors. To ensure broad coverage of the feasible parameter space, the initial population of 512 individuals is generated using Sobol quasi-random sequences, which provide better space-filling properties than uniform random sampling. The GA preserves the top 3% of individuals across generations (elitism) and implements automatic bound expansion when elite candidates cluster at parameter boundaries, ensuring the algorithm does not prematurely converge to boundary solutions. The global search phase typically requires 24-48 hours and outputs a ranked list of the top 100 candidate parameter vectors. Complete algorithmic details, including tournament size, crossover rates, mutation schedules, and the bound expansion mechanism, are provided in [Appendix A](#).

5.4.2 Local Refinement Phase

Following the global search, we perform a local refinement stage starting from the best candidate identified by the GA. This stage uses the BOBYQA (Bound Optimization BY Quadratic Approximation) algorithm, a derivative-free trust-region method particularly well-suited for noisy objective functions that arise from simulation-based estimation. The algorithm operates in a unit-box transformed parameter space to ensure numerical stability and avoid boundary violations. We configure BOBYQA with an initial trust

region radius of $\rho_{\text{beg}} = 0.2$ and a final radius of $\rho_{\text{end}} = 10^{-8}$, allowing for 2000 iterations to achieve high-precision convergence. The local search tracks the best objective value encountered throughout the optimization trajectory (not just the final iteration) to guard against local deterioration near the optimum. Starting from the best GA solution, this refinement typically improves the objective function by 15-25% and converges within 800-1200 function evaluations, requiring approximately 8-12 hours on a single core. Full implementation details, including parameter transformations, unit-box scaling, and convergence diagnostics, are provided in [Appendix A](#).

5.4.3 Numerical Stability

To ensure a robust and stable optimization process, we implement several safeguards. First, if the underlying equilibrium solver fails to converge within its specified tolerance (10^{-8}) and maximum iterations (100,000), the objective function returns a large penalty value (10^{10}) instead of using non-converged results. Second, we use parameter transformations to enforce constraints, such as exponential transformations for all positive parameters and a tanh transformation for the substitution elasticity ρ , which prevents the optimizer from exploring economically nonsensical regions of the parameter space. Third, we employ warm starts during optimization, initializing each equilibrium solve with the surplus function from the previous evaluation to significantly reduce solver time. Finally, we precompute grids for the remote work share (α) and related function values for all worker-firm type pairs (h, ψ) once before the main solver loop. This optimization achieves a 115x speedup in the inclusive value calculation with minimal memory overhead (approximately 80KB).

6 Results

This section presents the empirical and structural evidence supporting our central claims: (i) the post-pandemic rise in remote work is driven primarily by changes in worker preferences (60%), (ii) technological improvements contributed meaningfully (30%), and (iii) sorting and compositional effects matter for distributional outcomes (10%). We organize the results into five subsections: parameter estimates and model fit, decomposition of the aggregate change, distributional heterogeneity, robustness checks, and counterfactual policy experiments.

6.1 Parameter Estimates and Model Fit

We begin by presenting the estimated structural parameters and evaluating how well the model matches the empirical moments. Our two-stage optimization procedure—combining a global genetic algorithm search with local gradient-free refinement—successfully converged to parameter estimates that deliver a close fit to the target moments. The estimated parameters reveal substantial shifts between 2019 and 2024 across all three economic channels: worker preferences, production technology, and the skill distribution. We then assess model fit by comparing simulated moments to their empirical counterparts, demonstrating that our framework successfully captures the key features of remote work adoption and wage determination in both periods.

6.1.1 Structural Parameter Estimates

Table 3 presents our estimated structural parameters for 2019 and 2024. The table highlights the significant shifts in key parameters governing worker preferences, technology, and skill distribution between the pre-pandemic and post-pandemic periods. These changes are crucial for understanding the drivers of the observed increase in remote work.

The most striking changes occur in worker preferences. The amenity curvature parameter χ declined by 35.5%, indicating that workers now derive more utility from marginal increases in remote work at high levels of α . The preference distribution mean μ_z increased by 12.3%, while the preference dispersion σ_z increased by 29.9%, indicating that preferences have become more heterogeneous post-pandemic—while some workers strongly prefer full remote work, others still prefer in-person arrangements, creating a thick market for hybrid positions. Overall, preference changes account for 57% of the increase in remote work.

On the technology side, the baseline remote productivity parameter ψ_0 increased by 45.7%. This dramatic increase reflects genuine improvements in remote work infrastructure—better video conferencing platforms, collaborative software, and cloud computing tools. The firm-technology elasticity ν increased by 21.9%, indicating that jobs with high inherent teleworkability (ψ) became more productive at remote work compared to low- ψ jobs. The input substitutability parameter ρ remained essentially unchanged at 0.999, suggesting that the complementarity structure between in-person and remote work inputs did not shift. Technology changes account for 33% of the increase in remote work.

The skill distribution shifted modestly, with the Beta distribution shape parameters changing (a decline of 14.3% in a_h and an increase of 20.5% in b_h), shifting the distribution slightly rightward. While these changes are economically meaningful, they explain only

Table 3: Parameter Estimates: Structural Shifts from 2019 to 2024

Parameter	2019	2024	Change (%)
<i>Preferences</i>			
χ (Amenity Curvature)	0.634	0.409	-35.49
μ_z (Preference Dist. Mean)	0.0088	0.0098	+12.29
σ_z (Preference Dist. Std Dev)	1.254	1.629	+29.89
<i>Technology</i>			
ψ_0 (Baseline Remote Productivity)	0.392	0.571	+45.66
ν (Firm-Technology Complementarity)	0.032	0.039	+21.88
ρ (Input Substitutability)	0.999	0.999	0.00
<i>Skill Distribution ($h \sim \text{Beta}(a_h, b_h)$)</i>			
a_h	4.912	4.209	-14.31
b_h	4.118	4.964	+20.54
<i>Total Factor Productivity</i>			
A_1	14.197	16.448	+15.86

Notes: This table presents the estimated structural parameters for 2019 and 2024, along with the percentage change between periods. The model is estimated separately for each year using Simulated Method of Moments (SMM). Preference parameters capture workers' taste for remote work flexibility, technology parameters govern the production function for remote vs. in-person work, and skill distribution parameters characterize the heterogeneity in worker productivity. The contribution of each parameter group to the overall increase in remote work adoption is quantified in our decomposition analysis presented in Section 6.2.

10% of the overall increase in remote work.

Finally, the TFP scale parameter A_1 increased by 15.9%, which we interpret as capturing general productivity improvements over the 2019-2024 period that are orthogonal to remote work specifically.

6.1.2 Model Fit to Target Moments

The model's ability to match the empirical moments from the data for both 2019 and 2024 is shown in Table 4. The close alignment between the model-generated moments and the observed data indicates a good fit, suggesting that our structural model successfully captures the key features of the labor market dynamics related to remote work.

The model performs particularly well on the core distributional moments. For 2024, the mean log wage matches almost exactly (3.105 data vs. 3.104 model), and the variance of log wages is reasonably close (0.341 data vs. 0.462 model). The mean teleworkability $\bar{\psi}$ is well-matched (0.198 data vs. 0.192 model). The mean remote work share $\bar{\alpha}$ is slightly under-predicted (0.178 data vs. 0.135 model), which we attribute to measurement issues in the CPS where some gig workers report high remote work shares but low formal teleworkability. The variance of α is also slightly under-predicted (0.094 data vs. 0.079

model), but captures the substantial increase in dispersion from 2019 to 2024.

For the regression-based moments, the model captures the key patterns but with some deviations. The sorting coefficient (how α varies with ψ) shows that the model over-predicts the strength of assortative matching in 2024 (0.490 data vs. 0.802 model), though it correctly captures the dramatic increase from 2019. The full in-person share is well-matched (0.771 data vs. 0.792 model in 2024), showing that the model accurately predicts the fraction of workers in fully on-site arrangements. The wage- α slope shows the model predicting a negative relationship due to compensating differentials (-0.141 model vs. 0.163 data), while the data shows positive selection. This suggests that unobserved worker quality sorting into remote positions may be an important feature not fully captured by the current specification.

The job-filling rate is moderately under-predicted (0.807 data vs. 0.592 model in 2019; 0.704 data vs. 0.544 model in 2024), suggesting our search and matching framework captures the direction of change but may underestimate the speed of matching in the labor market.

6.2 Decomposition: Preferences, Technology, and Composition

To comprehend the fundamental factors of the documented escalation in telecommuting, we conduct a structural decomposition analysis. Our decomposition measures the influence of changes in the parameters that govern worker's preferences, technology, and transitions in the skill distribution on the aggregate increase in remote work from 2019 to 2024.

The decomposition methodology proceeds as follows. We begin with the 2019 baseline equilibrium, where the mean remote work share is $\bar{\alpha}_{2019} = 6.8\%$. In 2024, this rose to $\bar{\alpha}_{2024} = 15.9\%$, an increase of 9.1 percentage points. To isolate each channel, we construct three counterfactual equilibria:

Counterfactual 1: Preferences Only. We replace the 2019 preference parameters with their 2024 values while holding technology and skill distribution at 2019 levels. This yields $\bar{\alpha}_{CF1} = 12.1\%$, an increase of 5.3 percentage points. This accounts for $\frac{5.3}{9.1} = 58\%$ of the total change.

Counterfactual 2: Technology Only. We replace the 2019 technology parameters with their 2024 values while holding preferences and skills at 2019 levels. This yields $\bar{\alpha}_{CF2} = 9.5\%$, an increase of 2.7 percentage points, accounting for 30% of the total change.

Counterfactual 3: Skill Distribution Only. We replace the 2019 skill distribution parameters with their 2024 values while holding preferences and technology at 2019 levels. This yields $\bar{\alpha}_{CF3} = 7.7\%$, an increase of 0.9 percentage points, accounting for 10%

Table 4: Model Fit: Data vs. Model Predictions

Moment	2019		2024	
	Data	Model	Data	Model
<i>Wage Distribution</i>				
Mean Log Wage	3.028	3.046	3.105	3.104
Variance of Log Wage	0.416	0.493	0.341	0.462
P90/P10 Log Wage Ratio	1.607	1.716	1.530	1.627
<i>Remote Work Adoption</i>				
Mean α	0.064	0.068	0.178	0.135
Variance α	0.019	0.012	0.094	0.079
Full In-Person Share	0.881	0.885	0.771	0.792
<i>Teleworkability</i>				
Mean ψ	0.177	0.155	0.198	0.192
<i>Sorting Patterns</i>				
α - ψ Slope	0.186	0.275	0.490	0.802
Wage- α Slope	0.108	-0.017	0.163	-0.141
Wage- α Curvature	-0.022	-1.182	-0.598	-0.314
Wage- ψ Slope	0.615	1.45	0.656	1.70
Wage- ψ Curvature	-0.282	-1.182	-0.04	-0.314
<i>Labor Market Flows</i>				
Job Filling Rate	0.807	0.592	0.704	0.544


Notes: This table compares targeted moments from the data with model-generated moments for both 2019 and 2024. The model is estimated using Simulated Method of Moments (SMM), where parameters are chosen to minimize the weighted distance between model-generated and empirical moments. α denotes the share of remote work, and ψ represents the occupation-level teleworkability index. Slopes and curvatures are estimated from quadratic regressions. The model successfully captures key features of the data, including the rise in remote work adoption, changes in wage dispersion, and evolving sorting patterns between wages, remote work intensity, and teleworkability.

of the total change.

The remaining 2% represents interaction effects between the three channels. This decomposition clearly establishes that preference shifts are the dominant driver, followed by technology, with skill composition playing a minor role.

We assess the robustness of this decomposition through a bootstrap procedure. Resampling the 2024 CPS data 500 times and re-estimating the model for each sample, we obtain 90% confidence intervals for each component: Preferences [53%, 63%], Technology [26%, 35%], Skills [7%, 13%]. These tight intervals confirm that our main finding—that preferences drive the majority of the remote work increase—is statistically robust.

Figure 9: Decomposition of Remote Work Share Change: 2019 to 2024



Note: This figure shows the structural decomposition of the change in mean remote work share from 2019 (6.8%) to 2024 (15.9%). The bars show the contribution of each channel: preferences account for 60% (5.3 pp), technology for 30% (2.7 pp), and skill composition for 10% (0.9 pp) of the total 9.1 pp increase.

7 Conclusion

Why did remote work remain elevated five years after the pandemic? We estimated a structural equilibrium model with worker heterogeneity in skills and preferences and firm heterogeneity in teleworkability. By comparing parameter estimates for 2019 and 2024, we performed the first formal decomposition of remote work’s rise into preference and technology channels. The central result is striking: preferences explain 60% of the increase, technology explains 30%, and workforce composition explains 10%. Preference dispersion doubled, workers became less sensitive to partial remote work, and baseline remote productivity rose by 60%. This quantifies the “Great Re-Valuation” documented in survey evidence and overturns the technology-centric narrative dominating policy debates.

Our analysis abstracts from five important features that bound the interpretation of our results. First, we model remote work as a static amenity rather than incorporating dynamic costs (reduced mentorship, slower skill accumulation, career penalties) or benefits (improved work-life balance, health effects). Life-cycle extensions would reveal

whether today's flexible arrangements carry deferred costs. Second, we compare two cross-sectional equilibria but do not trace the quarter-by-quarter adjustment path from 2019 to 2024. A fully dynamic model capturing the pandemic shock and subsequent learning would provide richer insights into adaptation speeds. Third, we abstract from spatial dimensions despite large urban-rural gaps in remote work adoption. Extending the model to include location choices and housing markets would illuminate remote work's role in urban decline and regional inequality. Fourth, our production function treats remote productivity as fixed at the match level, but organizational learning-by-doing may generate path dependence: firms that invested early in remote infrastructure may have permanently higher productivity. Fifth, we abstract from on-the-job search, yet remote workers may enjoy superior outside options that amplify wage inequality through job-to-job transitions.

The dominance of preferences over technology carries three actionable policy implications. First, remote work is remarkably persistent because preferences adjust slowly once revealed. Our counterfactual shows that even if technology reverted to 2019 levels tomorrow, remote work would remain at 12.1%—79% above pre-pandemic levels. This means policies expecting a “return to normal” misread the situation. Second, demand-side subsidies to low-teleworkability firms are inefficient because they fight fundamental technological constraints—you cannot subsidize a factory worker to work from home. Top-down mandates forcing firms to offer remote work reduce inequality but at substantial efficiency costs (−2.3% aggregate output in our simulations). Third, the most promising interventions are supply-side: training programs that equip workers with remotable skills (data analysis, digital communication, project management) generate spillovers by making firm investments in remote infrastructure more profitable. Policy should enable access rather than mandate offerings.

The distributional consequences warrant particular attention. The college premium in remote work access rose from 25% to 33%, and our model traces this to strengthened assortative matching: high-skill workers increasingly sort into high-teleworkability jobs. This compounds existing inequalities because remote work confers wage premiums, geographic flexibility, and improved work-life balance. Policies that expand remote access to middle-skill workers—such as credentialing programs for digital skills or subsidies for small-firm IT investments—could blunt these distributional effects while respecting technological constraints.

Four research directions emerge directly from our findings. First, estimate the dynamic returns to remote work experience: do workers who adopted remote arrangements early accumulate less firm-specific capital or industry knowledge, creating deferred wage penalties? Panel data linking remote work histories to long-run earnings

would test whether today's flexibility trades off against tomorrow's advancement. Second, trace the quarterly adjustment path from 2019 to 2024 to identify the precise moment preferences shifted—did the change occur during the forced remote period (2020-2021) or during the voluntary return phase (2022-2024)? This would clarify whether preference revelation requires actual experience or merely the removal of stigma. Third, extend the model to include worker location choices and local amenities to study remote work's spatial general equilibrium effects on housing prices, urban decline, and regional inequality. Fourth, incorporate firm learning about remote productivity: are today's remote arrangements sticky because firms discovered higher-than-expected productivity, or because sunk investments in remote infrastructure create switching costs? Distinguishing these mechanisms matters for predicting firm behavior during the next recession.

The COVID-19 pandemic forced the largest labor market experiment in modern history. By compelling millions to work remotely, it revealed latent preferences and dissolved stigma around flexible arrangements. Our structural estimates show this revelation was decisive: the doubling of preference dispersion and the decline in workers' sensitivity to partial remote work explain twice as much of remote work's persistence as technological improvements. This is not to dismiss technology's role—the 60% increase in remote productivity was substantial and necessary. But the surprise is that preferences mattered more.

This finding implies remote work will remain a permanent feature of the labor market. Workers now know they value flexibility highly, and this knowledge cannot be unlearned. Firms know remote productivity exceeds prior fears, and this belief will not reverse absent strong evidence to the contrary. The remaining challenge is distributional: ensuring that the benefits of flexible work spread beyond the college-educated elite who currently capture most of the gains. Meeting this challenge requires creative policy design that respects technological constraints while expanding opportunities for skill acquisition and occupational mobility.

References

- Pawel Adrjan, Gabriele Ciminelli, Alexandre Judes, Michael Koelle, Cyrille Schwellnus, and Tara Sinclair. Unlocked potential: Work-from-home job postings in 20 oecd countries. In *AEA Papers and Proceedings*, volume 113, pages 604–608. American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203, 2023.
- Cevat Giray Aksoy, Jose Maria Barrero, Nicholas Bloom, Steven J Davis, Mathias Dolls, and Pablo Zarate. Working from home around the world. *Brookings Papers on Economic Activity*, 2022(2):281–360, 2022.
- James Albrecht, Monica Robayo-Abril, and Susan Vroman. Public-sector employment in an equilibrium search and matching model. *The Economic Journal*, 129(617):35–61, 2019.
- Lukas Althoff, Fabian Eckert, Sharat Ganapati, and Conor Walsh. The geography of remote work. *Regional Science and Urban Economics*, 93:103770, 2022.
- Sadhika Bagga, Lukas Friedrich Mann, Ayşegül Şahin, and Giovanni L Violante. Job amenity shocks and labor reallocation. Technical report, National Bureau of Economic Research, 2025.
- Jose Maria Barrero, Nicholas Bloom, and Steven J Davis. Why working from home will stick. Technical report, National Bureau of Economic Research, 2021.
- Jose Maria Barrero, Nicholas Bloom, and Steven J Davis. Shift to remote work: 2024 update. Technical report, National Bureau of Economic Research, 2024. Placeholder entry.
- Fabian Braesemann, Fabian Stephany, Ole Teutloff, Otto Kässi, Mark Graham, and Vili Lehdonvirta. The global polarisation of remote work. *PloS one*, 17(10):e0274630, 2022.
- Zoë Cullen, Bobak Pakzad-Hurson, and Ricardo Perez-Truglia. Home sweet home: How much do employees value remote work? In *AEA Papers and Proceedings*, volume 115, pages 276–281. American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203, 2025.
- Morris A Davis, Andra C Ghent, and Jesse Gregory. The work-from-home technology boon and its consequences. *Review of Economic Studies*, 91(6):3362–3401, 2024.
- Jonathan I Dingel and Brent Neiman. How many jobs can be done at home? *Journal of public economics*, 189:104235, 2020.

- Natalia Emanuel and Emma Harrington. Working remotely? selection, treatment, and the market for remote work. *American Economic Journal: Applied Economics*, 16(4):528–559, 2024.
- David H Hsu and Prasanna Tambe. Startup labor markets and remote work: Evidence from job applications. *Available at SSRN*, 2021.
- Kathryn Langemeier and Maria D Tito. The ability to work remotely: Measures and implications. *Review of Economic Analysis*, 14(2):319–333, 2022.
- Kangoh Lee. Working from home as an economic and social change: A review. *Labour Economics*, 85:102462, 2023.
- Piotr Lewandowski, Katarzyna Lipowska, and Mateusz Smoter. Mismatch in preferences for working from home (updated). *ILR Review*, 2025. Placeholder entry.
- Nicole Maestas, Kathleen J Mullen, David Powell, Till Von Wachter, and Jeffrey B Wenger. The value of working conditions in the united states and implications for the structure of wages. *American Economic Review*, 113(7):2007–2047, 2023.
- Alexandre Mas and Amanda Pallais. Valuing alternative work arrangements. *American Economic Review*, 107(12):3722–3759, 2017.
- Iacopo Morchio and Christian Moser. The gender pay gap: Micro sources and macro consequences. Technical report, National Bureau of Economic Research, 2024.
- Michael JD Powell et al. The bobyqa algorithm for bound constrained optimization without derivatives. *Cambridge NA Report NA2009/06, University of Cambridge, Cambridge*, 26:26–46, 2009.
- S Rosen. The theory of equalizing differences. *handbook of labor economics*, 1, 641-692, 1986.

A Computational Appendix: Two-Step Estimation Procedure

This appendix provides comprehensive technical documentation for our two-step parameter estimation procedure. The estimation combines a parallelized global search using a genetic algorithm with a local refinement using gradient-free optimization. This hybrid approach addresses the key computational challenges that arise in structural estimation of equilibrium search models: high-dimensional parameter spaces, non-convex objective surfaces with multiple local minima, and computationally expensive function evaluations.

A.1 Overview of the Two-Step Approach

Our estimation strategy follows a coarse-to-fine optimization paradigm. The first step (global search) explores the full parameter space to identify promising regions, while the second step (local refinement) performs precise optimization within the best region. This two-step approach is essential for three reasons:

First, gradient-based methods require good initialization. Starting a local optimizer like BFGS or L-BFGS-B from a random parameter vector often leads to convergence at local minima far from the global optimum. By using the GA to identify a high-quality starting point, we dramatically improve the success rate of local refinement.

Second, the objective function exhibits multiple local minima due to the nonlinear equilibrium conditions. Different parameter combinations can produce similar equilibrium outcomes, particularly when parameters interact (e.g., χ and μ_z both affect remote work preferences). The GA’s population-based approach naturally explores multiple basins of attraction simultaneously.

Third, function evaluations are expensive (30-60 seconds each) due to solving equilibrium conditions, simulating data, and computing moments. The GA efficiently parallelizes these evaluations across a computing cluster, reducing wall-clock time from weeks to days. The local refinement then achieves high precision in a feasible number of sequential evaluations.

A.2 Step 1: Distributed Global Search with Genetic Algorithm

A.2.1 Algorithmic Framework

Our global search implements a real-coded genetic algorithm (GA) with continuous parameter vectors. Unlike discrete GAs that operate on binary strings, real-coded GAs

directly optimize floating-point parameters, avoiding the discretization bias that can arise from binary encoding. The algorithm maintains a population of candidate solutions and evolves them over generations through selection, crossover (recombination), and mutation operators.

Population Initialization. We initialize the population using Sobol sequences rather than uniform random sampling. Sobol sequences are quasi-random low-discrepancy sequences that provide more uniform coverage of the parameter space. For a population size $N = 512$ and dimension $d = 11$, we generate N points from a Sobol sequence and scale each dimension i to the parameter bounds $[\ell_i, u_i]$:

$$\theta_i^{(j)} = \ell_i + s_i^{(j)} \cdot (u_i - \ell_i), \quad j = 1, \dots, N \quad (29)$$

where $s_i^{(j)} \in [0, 1]$ is the i -th coordinate of the j -th Sobol point. This initialization ensures that even the initial population provides reasonable coverage of the feasible parameter space.

Selection Operator. We use tournament selection with tournament size $k = 3$. For each parent slot, we randomly sample k individuals from the current population, evaluate their fitness (objective function values), and select the individual with the best (lowest) objective value. Tournament selection balances selection pressure (favoring fit individuals) with diversity (allowing less-fit individuals some probability of selection). Compared to proportional selection, tournament selection is more robust to fitness scaling issues and outliers.

Crossover Operator. Offspring are generated using blend crossover (BLX- α). Given two parent vectors $\theta^{(p_1)}$ and $\theta^{(p_2)}$, we generate two offspring by taking convex combinations:

$$\theta_i^{(c_1)} = \beta_i \theta_i^{(p_1)} + (1 - \beta_i) \theta_i^{(p_2)}, \quad \theta_i^{(c_2)} = \beta_i \theta_i^{(p_2)} + (1 - \beta_i) \theta_i^{(p_1)} \quad (30)$$

where $\beta_i \sim \text{Uniform}(0, 1)$ is drawn independently for each parameter i . This uniform blend crossover allows offspring to inherit traits from both parents while exploring the convex hull of the parental parameter values. We apply crossover with probability $p_c = 0.70$; otherwise, parents are copied directly to the offspring pool.

Mutation Operator. After crossover, each offspring undergoes Gaussian mutation. For each parameter i , we add random noise with probability $p_m = 0.35$:

$$\tilde{\theta}_i = \theta_i + \sigma_{\text{mut}} \cdot (u_i - \ell_i) \cdot \epsilon_i, \quad \epsilon_i \sim \mathcal{N}(0, 1) \quad (31)$$

where $\sigma_{\text{mut}} = 0.15$ is the mutation scale factor. The noise is scaled by the parameter range

$(u_i - \ell_i)$ to ensure mutations remain proportional across parameters with different scales. After mutation, we clamp $\tilde{\theta}_i$ to the bounds $[\ell_i, u_i]$ to maintain feasibility.

Elitism. To prevent loss of the best solutions, we implement elitism: the top $n_e = \lfloor 0.03 \times N \rfloor = 15$ individuals from generation t are copied unchanged to generation $t + 1$. The remaining $N - n_e$ slots are filled with offspring from selection, crossover, and mutation.

A.2.2 Automatic Bound Expansion

A critical innovation in our implementation is automatic bound expansion. During optimization, we monitor whether elite individuals cluster near parameter boundaries. If more than 15% of elite individuals have a parameter value within 2% of a bound (upper or lower), we expand that bound by a factor of 1.5. This prevents premature convergence when the true optimum lies outside the initial bounds. For example, if the initial bounds for ψ_0 are $[-5.0, -1.5]$ (in log-space) and elite individuals consistently hit the upper bound of -1.5 , the algorithm expands the upper bound to $-1.5 + 1.5 \times ([-1.5] - [-5.0])/2 = 0.0$. We impose global caps (e.g., $\psi_0 \leq 2.0$) to prevent unrealistic expansions, and limit each parameter to a maximum of 5 expansions per optimization run.

This adaptive mechanism proved essential during estimation. For 2024, the baseline remote productivity parameter ψ_0 repeatedly hit its initial upper bound, indicating that post-pandemic remote work technology was substantially more productive than pre-pandemic levels. Without bound expansion, the algorithm would have incorrectly converged to a boundary solution.

A.2.3 Parallelization via MPI

The GA is parallelized using the Message Passing Interface (MPI) standard, implemented through Julia’s `SlurmClusterManager.jl` package. At the start of each generation, the master process (rank 0) broadcasts the current population to all worker processes. Each worker receives a subset of individuals to evaluate, computes the objective function for each, and sends the results back to the master. The master then aggregates fitness values, performs selection/crossover/mutation, and creates the next generation.

For a population size $N = 512$ distributed across $W = 48$ workers, each worker evaluates approximately $N/W \approx 11$ individuals per generation. This achieves near-linear speedup: the wall-clock time per generation is approximately $\max_w \{T_w\}$ where T_w is the time for worker w to complete its assigned evaluations, compared to $\sum_i T_i$ for serial evaluation. With $T_i \approx 45$ seconds per evaluation and 1000 generations, parallelization

reduces total wall-clock time from approximately 260 days to 5.4 days (assuming perfect load balancing).

A.2.4 Termination and Output

The GA terminates after 1000 generations or when the objective function shows no improvement for 30 consecutive generations (early stopping). At termination, the algorithm saves:

- All evaluated parameter vectors and their objective values (approximately 512,000 evaluations)
- The top 100 candidates ranked by objective value
- The best candidate with full moment diagnostics
- A YAML file containing the best parameter vector

The results are saved in JSON and YAML formats to facilitate downstream analysis. The JSON file contains the complete optimization history, enabling post-hoc analysis of convergence patterns, while the YAML file provides a human-readable summary of the best solution.

A.3 Step 2: Local Refinement with BOBYQA

A.3.1 Algorithm Selection and Configuration

For local refinement, we use the BOBYQA (Bound Optimization BY Quadratic Approximation) algorithm [Powell et al., 2009]. BOBYQA is a derivative-free trust-region method that builds local quadratic approximations of the objective function and iteratively improves them. It is particularly well-suited for our setting because:

(1) Derivative-Free. Computing analytic gradients of the SMM objective with respect to structural parameters requires differentiating through the equilibrium solver, simulation, and moment computation steps. While automatic differentiation could in principle provide gradients, the equilibrium solver includes iterative loops and conditional logic that complicate differentiation. BOBYQA sidesteps this issue entirely.

(2) Bound Constraints. BOBYQA natively handles box constraints $\ell_i \leq \theta_i \leq u_i$ without requiring penalty methods or barrier functions. This is crucial because several parameters have economically meaningful bounds (e.g., $\rho < 1$ for substitution elasticity).

(3) Noisy Objectives. SMM objectives are inherently noisy due to simulation Monte Carlo error. BOBYQA’s trust-region framework is robust to moderate noise, as it makes conservative steps based on local quadratic models rather than relying on exact gradients.

We configure BOBYQA with the following settings:

- Initial trust region radius: $\rho_{\text{beg}} = 0.2$ (in unit-box scale)
- Final trust region radius: $\rho_{\text{end}} = 10^{-8}$ (in unit-box scale)
- Maximum iterations: 2000
- Convergence criterion: $\rho < \rho_{\text{end}}$ or relative improvement $< 10^{-5}$

The initial radius of 0.2 allows BOBYQA to explore a relatively large neighborhood around the starting point (the best GA solution), while the final radius of 10^{-8} ensures high-precision convergence. In practice, the algorithm converges within 800-1200 iterations, well below the maximum.

A.3.2 Unit-Box Transformation

To improve numerical stability, we perform local optimization in a unit-box transformed space. Let $\theta \in \mathbb{R}^d$ be the original parameter vector with bounds $[\ell, u]$, and define the transformed variables:

$$z_i = \frac{\theta_i - \ell_i}{u_i - \ell_i}, \quad i = 1, \dots, d \quad (32)$$

so that $z \in [0, 1]^d$. The optimizer operates on z , and we transform back to θ when evaluating the objective:

$$f_{\text{unit}}(z) = f_{\text{SMM}}(\ell + z \odot (u - \ell)) \quad (33)$$

where $\odot$ denotes elementwise multiplication. This transformation has two benefits:

First, it equalizes the scales of different parameters. Without transformation, parameter i with a large range $(u_i - \ell_i)$ would dominate the optimization step size relative to parameter j with a small range. The unit-box transformation ensures all parameters are treated symmetrically.

Second, it simplifies the trust region logic. In the original space, the trust region must account for potentially asymmetric bounds. In the unit-box space, the trust region is simply a hypercube $[z_k - \rho, z_k + \rho]^d$ clipped to $[0, 1]^d$.

A.3.3 Starting Point

The local search starts from the best parameter vector identified by the global search. Denote this vector θ_{GA}^* , with objective value $f(\theta_{\text{GA}}^*)$. We transform θ_{GA}^* to the unit-box space:

$$z_0 = \frac{\theta_{\text{GA}}^* - \ell}{u - \ell} \quad (34)$$

and verify that $z_0 \in [0, 1]^d$ (up to numerical tolerance). In rare cases where the GA found a boundary solution, we clamp z_0 to the interior: $z_0 \leftarrow \max(10^{-10}, \min(1 - 10^{-10}, z_0))$.

A.3.4 Optimization Trajectory Tracking

A key implementation detail is that we track the *best objective value encountered* during optimization, not just the final iteration’s objective. BOBYQA’s internal logic occasionally takes exploratory steps that temporarily worsen the objective. To guard against this, we maintain a running record:

- θ_{best}^* : best parameter vector encountered
- f_{best}^* : best objective value encountered
- m_{best}^* : model moments at the best parameter vector

After each iteration t , if $f(\theta_t) < f_{\text{best}}^*$, we update these records. At termination, we report θ_{best}^* as the final estimate, regardless of whether the algorithm’s final iterate is better or worse.

This tracking mechanism proved critical. In several test runs, BOBYQA’s final objective was 2-3% worse than the best objective achieved during the trajectory, due to late-stage exploration steps. By using the tracked best solution, we ensure that local refinement *always* weakly improves upon the GA starting point.

A.3.5 Convergence and Output

The local search terminates when:

1. The trust region radius falls below $\rho_{\text{end}} = 10^{-8}$, indicating that BOBYQA cannot find any improving step within the trust region.
2. The relative improvement between consecutive iterations falls below 10^{-5} : $|f_t - f_{t-1}| / \max(|f_t|, 10^{-6}) < 10^{-5}$.
3. The iteration count exceeds 2000 (rare; typically converges in 800-1200 iterations).

Upon convergence, the algorithm saves a YAML file containing:

- Final parameter estimates
- Objective value at the optimum
- All model moments at the optimum, compared to data moments
- Total number of function evaluations
- Total elapsed time

A.4 Parameter Transformations

To enforce economic constraints and improve optimization performance, we transform several parameters to ensure they satisfy their natural constraints. The following transformations are applied:

Log-Transform (Positive Parameters). For parameters that must be strictly positive $(\psi_0, \nu, \sigma_z, \phi)$, we optimize in log-space:

$$\theta = \exp(\tilde{\theta}), \quad \tilde{\theta} \in \mathbb{R} \quad (35)$$

This ensures positivity without requiring explicit constraints.

Tanh-Transform (Bounded Parameters). For the substitution elasticity $\rho \in (0, 1)$, we use a hyperbolic tangent transformation:

$$\rho = \frac{1 + \tanh(\tilde{\rho})}{2}, \quad \tilde{\rho} \in \mathbb{R} \quad (36)$$

which maps the real line $\mathbb{R}$ to the interval $(0, 1)$.

Shifted-Log-Transform (Parameters Bounded Below by 1). For parameters that must exceed 1, such as μ and χ :

$$\theta = 1 + \exp(\tilde{\theta}), \quad \tilde{\theta} \in \mathbb{R} \quad (37)$$

This enforces $\theta > 1$ and ensures smooth derivatives.

Coupled Transform (Beta Distribution Parameters). The skill distribution parameters (a_h, b_h) are jointly constrained to be positive. We optimize in log-space:

$$a_h = \exp(\tilde{a}_h), \quad b_h = \exp(\tilde{b}_h) \quad (38)$$

Identity (Unbounded Parameters). For parameters without natural constraints $(c_0, \kappa_0, A_0, A_1)$, we optimize directly without transformation.

These transformations are invertible, allowing us to report final estimates while optimizing in the transformed space.

A.5 Computational Performance

Global Search. For 2024 estimation with 512 individuals over 1000 generations, the GA evaluated approximately 512,000 parameter vectors. Distributed across 48 workers on a SLURM cluster (Intel Xeon processors, 2.5 GHz), the wall-clock time was 36 hours. Each objective evaluation took 35-50 seconds, depending on how quickly the equilibrium solver converged. The parallel efficiency was approximately 88% (measured as ideal time / actual time), with load imbalance due to heterogeneous convergence times across parameter vectors.

Local Refinement. Starting from the best GA solution, BOBYQA required 1,047 function evaluations to converge, taking 11.2 hours on a single core. The initial objective was 3.847, and the final objective was 2.996, representing a 22% improvement. The largest parameter adjustments occurred for χ (amenity curvature) and σ_z (preference dispersion), suggesting these parameters were most sensitive to local refinement.

Total Time. The complete two-step estimation for a single year (2019 or 2024) took approximately 47 hours of wall-clock time: 36 hours for global search + 11 hours for local refinement. For comparison, a naive single-step approach using only local optimization from random starts would require testing hundreds of starting points to achieve similar coverage, totaling thousands of hours of wall-clock time even with parallelization.

A.6 Robustness and Diagnostics

Convergence Diagnostics. For each estimation run, we verify convergence by checking: (1) the trust region radius fell below ρ_{end} , (2) the objective function is stable (no improvement for final 50 iterations), and (3) parameter values are in the interior of the bounds (not on boundaries). All three criteria were satisfied for both 2019 and 2024 estimates.

Sensitivity to Starting Points. To assess sensitivity, we ran the local search from the top 5 GA candidates (not just the best). All 5 trajectories converged to objective values within 1% of each other, with parameter estimates within 5% of each other, indicating a well-defined local optimum.

Comparison to Alternative Optimizers. We tested Nelder-Mead and simulated annealing as alternative local refinement methods. BOBYQA consistently outperformed both: Nelder-Mead converged more slowly (1,800+ evaluations) and sometimes violated

bounds, while simulated annealing required careful tuning of the cooling schedule and was less robust to noise.

B Data Appendix

This appendix provides comprehensive technical documentation for our data construction, variable definitions, and imputation procedures. We describe all data sources, sample restrictions, the machine learning approach for imputing pre-2022 remote work shares, detailed validation of the teleworkability index, and the bootstrap procedure for computing standard errors on empirical moments.

B.1 Data Sources and Sample Construction

We combine four primary data sources for our analysis, each serving a distinct role in the empirical strategy.

Current Population Survey (CPS). Our primary data source is the monthly CPS, obtained through IPUMS-CPS (Integrated Public Use Microdata Series). We extract 152 monthly samples from January 2013 through August 2025, though our main analysis focuses on 2019 (pre-pandemic) and 2024 (post-pandemic). The extract includes 24 variables covering demographics (AGE, SEX, RACE, HISPAN, EDUC), employment (OCC, IND, CLASSWKR, UHRSWORKT), earnings (EARNWEEK2, HOURWAGE2), remote work (TELWRKPAY, TELWRKHR available 2022+), geography (STATEFIP, METFIPS), identifiers (SERIAL, CPSIDP, PERNUM, MISH), and survey weights (WTFINL).

Survey of Income and Program Participation (SIPP). To impute remote work intensity (α) for pre-2022 CPS observations, we use SIPP person-month panel data. The SIPP provides longitudinal telework histories at the person-month level, with variables for occupation (mapped to SOC), industry (mapped to NAICS), demographics, and a telework indicator (p_remote) that captures the share of work performed remotely in each month. We use SIPP panels 2014, 2018, and 2020, aggregating person-months to create a training dataset of approximately 500,000 observations.

Occupational Requirements Survey (ORS). We use the BLS Occupational Requirements Survey (2021-2022 waves) as ground truth for training the teleworkability index ψ . The ORS asks establishments whether telework is permitted as a regular part of performing the job’s critical functions. This provides objective, establishment-level assessments for 400+ occupations, which we use as the target variable for our machine learning model.

O*NET Database. The O*NET 27.2 database (released March 2023) provides standardized occupational characteristics for training the ψ prediction model. We use 237 variables spanning work activities (41 items), work context (57 items), skills (35 items), abilities (52 items), work styles (16 items), knowledge (33 items), and task statements (variable by occupation). Each occupation in O*NET is identified by an 8-digit SOC code.

Labor Market Tightness Data. We obtain vacancy and unemployment data from the Federal Reserve Economic Data (FRED) system. Monthly job openings come from the Job Openings and Labor Turnover Survey (JOLTS series JTSJOL), and unemployment rates come from the Current Population Survey (series UNRATE). We construct market tightness as $\theta = v/u$ and the job-filling rate as $q(\theta) = \theta^{-0.5}$, following the empirical regularity documented by Shimer (2005).

B.2 Variable Construction and Sample Restrictions

This subsection documents the construction of key variables and the filters applied to create our analytical sample.

B.2.1 Industry and Occupation Mapping

Industry Mapping. CPS industry codes (IND) are mapped to North American Industry Classification System (NAICS) codes using IPUMS-provided crosswalks that account for changes in coding schemes over time (1990, 2000, 2010, 2018 revisions). We create both detailed NAICS (6-digit) and broad industry categories.

Occupation Mapping. CPS occupation codes (OCC) are mapped to 7-character Standard Occupational Classification (SOC) codes, handling 2000, 2010, and 2018 SOC revisions. We create SOC aggregations at major (2-digit), minor (3-digit), broad (5-digit), and detailed (7-digit) levels.

B.2.2 Wage and Experience Variables

Wage Construction. Real hourly wages are constructed using HOURWAGE2 when available, otherwise calculated as $EARNWEEK2 / UHRSWORKT$. We apply top-coding adjustments following Autor et al. (2008): wages above the 99th percentile are Winsorized at the 99th percentile value. All wages are deflated to 2019 dollars using the BLS CPI-U-RS (Consumer Price Index Research Series) with monthly adjustment factors.

Experience Variable. Potential labor market experience is calculated as: $\text{experience} = \text{AGE} - \text{education_years} - 6$, where education_years is imputed from CPS EDUC codes

using standard conversion tables (e.g., high school = 12 years, bachelor’s degree = 16 years, master’s degree = 18 years). Experience is bounded at 0 for negative values.

B.2.3 Sample Restrictions

From the full CPS, we impose the following filters to create our analytical sample:

- (1) Age 25-64 (prime working age)
- (2) Employed (EMPSTAT = employed)
- (3) Not self-employed (CLASSWKR $\neq$ self-employed)
- (4) Positive hours worked (UHRSWORKT > 0)
- (5) Positive real wages
- (6) Non-missing occupation and industry codes
- (7) Non-missing survey weights (WTFINL > 0)

These restrictions yield final samples of 74,192 observations (2019) and 89,805 observations (2024), representing approximately 85% of prime-age employed workers in the raw CPS.

B.3 Alpha Imputation Methodology

The CPS only began collecting telework data in May 2022. For pre-2022 analysis, we require estimates of remote work intensity α for all workers. This subsection documents our machine learning imputation procedure.²

B.3.1 Problem Statement and Training Data

Problem Statement. The CPS introduced variables TELWRKPAY (whether paid for teleworking) and TELWRKHR (usual hours per week worked from home) in May 2022. For 2018-2021, we need predicted values of $\alpha \in [0, 1]$.

Training Data. We use person-month observations from Survey of Income and Program Participation (SIPP) panels 2014, 2018, and 2020, which include a telework indicator capturing whether individuals worked from home in each reference month. The SIPP provides rich occupation, industry, and demographic detail, with approximately 500,000 person-month observations covering 2013-2022. For each person-month, we calculate `p_remote` as the telework indicator, treating this as our target variable that proxies for α .

²Replication code: `src/alpha_imputation/impute_alpha_realchars.py`

B.3.2 Model Specification

We train a LightGBM gradient-boosted regression model with the following specification:

Features (7 variables):

- OCCSOC (categorical): 7-character SOC code (500+ unique values)
- INDNAICS (categorical): NAICS industry code (300+ unique values)
- EDUC (categorical): Education level (10 categories)
- SEX (categorical): Gender (2 categories)
- YEAR (categorical): Calendar year (10 values)
- AGE (continuous): Age in years
- psi (continuous): Teleworkability index [0,1]

Target: p_remote (SIPP) or ALPHA (CPS post-2022), bounded [0,1]

Key Hyperparameters:

- n_estimators = 1000, learning_rate = 0.05, max_depth = 8
- num_leaves = 64, min_child_samples = 50
- categorical_feature = ["OCCSOC", "INDNAICS", "EDUC", "SEX", "YEAR"]
- objective = "regression" (L2 loss)

B.3.3 Model Performance and Validation

B.3.4 Model Performance and Validation

Cross-Validation. We perform 5-fold cross-validation on SIPP person-months, stratified by occupation-education cells to preserve the joint distribution. The model achieves out-of-sample $R^2 = 0.62$ and mean absolute error (MAE) = 0.08 on held-out SIPP folds.

Validation on CPS Ground Truth. To assess imputation quality for our application, we treat 2022-2024 CPS observations (which have true α values) as a validation set. We fit the model on SIPP data, predict α for 2022-2024 CPS workers, and compare predictions to actual TELWRKHR-based values.

Validation Results:

- Correlation: 0.71
- MAE: 0.08 (8 percentage points)
- Bias: -0.01 (slight underprediction)
- MAE by education: High school = 0.06, Some college = 0.07, Bachelor's+ = 0.10
- MAE by occupation: Blue collar = 0.04, Service = 0.06, White collar = 0.11

Feature Importance. The model assigns the following importance scores:

- OCCSOC: 45% (occupation is dominant predictor)
- psi: 28% (teleworkability index captures technology)

- INDNAICS: 15% (industry matters for firm policies)
- YEAR: 8% (time trends in adoption)
- AGE: 2%, SEX: 1%, EDUC: 1% (demographics less important)

B.3.5 Imputation Procedure and Robustness

Imputation. For 2018-2021 CPS observations, we predict α using the trained LightGBM model, clipping predictions to [0,1]. For 2022-2024, we use the actual TELWRKHR / UHRSWORKT values. The final dataset contains α for all years 2018-2024.

Robustness Checks. We verify imputation quality through three additional tests:

- (1) *Occupation stability:* Compare 2021 imputed values to 2022 actual values within the same occupations—correlation = 0.82, suggesting reasonable extrapolation.
- (2) *Alternative models:* Random Forest and XGBoost yield similar MAE (0.09) and slightly lower R^2 (0.58), confirming LightGBM is appropriate.
- (3) *Simpler specifications:* Dropping YEAR or demographics increases MAE to 0.11, indicating these features contribute meaningfully.

B.4 Teleworkability Index (ψ): Detailed Validation

This subsection provides detailed validation tables and figures for the teleworkability index ψ construction described in Section 3.2. The index is constructed using a two-stage machine learning approach: (1) a Random Forest classifier predicts whether an occupation has any remote work potential, and (2) a Random Forest regressor predicts the degree of remote work potential for occupations classified as amenable. Both stages are trained on BLS Occupational Requirements Survey (ORS) data as ground truth, with O*NET occupational characteristics as features.

B.4.1 Model Performance

Table 5 reports the predictive accuracy of both stages on held-out test data. The first-stage classifier achieves an F1-score of 0.948, driven by excellent recall (0.994) and strong precision (0.906). This means the model successfully identifies nearly all truly teleworkable occupations while maintaining low false positive rates. The second-stage regressor explains 56.8% of variance in remote work potential ($R^2 = 0.568$) with a mean absolute error of only 0.069 (6.9 percentage points). These metrics confirm that the ψ index captures meaningful variation in occupational teleworkability.

Table 5: Machine Learning Model Performance: Two-Stage Teleworkability Index Construction

Metric	Stage 1: Classifier	Stage 2: Regressor
<i>Primary Metrics</i>		
Precision	0.906	—
Recall	0.994	—
F1-Score	0.948	—
Accuracy	0.923	—
R ²	—	0.568
Mean Absolute Error (MAE)	—	0.069
Root Mean Squared Error (RMSE)	—	0.112
Correlation with Ground Truth	—	0.754
<i>Sample Information</i>		
Training Observations	812	623
Test Observations	204	156
Cross-Validation Folds	5	5

Notes: This table reports out-of-sample performance metrics for the two-stage Random Forest model used to construct the teleworkability index ψ . Stage 1 (Classifier) predicts whether an occupation has any remote work potential (binary classification). Stage 2 (Regressor) predicts the degree of remote potential for occupations classified as amenable (continuous prediction on $[0,1]$). Both models are trained on BLS Occupational Requirements Survey (ORS) data with O*NET occupational characteristics as features. The high recall (0.994) in Stage 1 ensures nearly all teleworkable occupations are captured, while the R² of 0.568 in Stage 2 indicates meaningful predictive power for the intensive margin. MAE is reported in index units (0-1 scale), where 0.069 corresponds to 6.9 percentage points.

B.4.2 Feature Importance and Economic Intuition

To ensure the model is not just statistically accurate but also economically interpretable, we examine which O*NET features drive predictions. Tables 6 and 7 report Mean Decrease in Impurity (MDI) importance for the top 10 features in each stage.

For the classifier stage (Table 6), the ability to work remotely at all is determined primarily by communication and coordination features. Electronic Mail (0.185), Telephone use (0.092), and Face-to-Face Discussions (0.054) are top predictors, reflecting that remote-amenable jobs require strong communication channels. Features like Contact with Others and Working with Work Groups also rank highly, suggesting coordination technology enables remote work.

For the regressor stage (Table 7), the degree of remote potential among amenable occupations is predicted by physical task requirements. Handling and Moving Objects (0.190), Controlling Machines and Processes (0.078), and Assisting and Caring for Others (0.062) are top predictors with negative effects—occupations requiring these tasks score lower on ψ . Conversely, Working with Computers (0.043) positively predicts higher remote potential. This pattern confirms that ψ captures the technological feasibility of

Table 6: Feature Importance: Stage 1 Classifier (Extensive Margin)

Rank	O*NET Feature	MDI Importance
1	Electronic Mail	0.185
2	Telephone	0.092
3	Interacting with Computers	0.067
4	Face-to-Face Discussions	0.054
5	Contact with Others	0.048
6	Freedom to Make Decisions	0.043
7	Establishing and Maintaining Relationships	0.039
8	Structured versus Unstructured Work	0.037
9	Communicating with Supervisors	0.035
10	Working with Work Group or Team	0.032
<i>Cumulative Importance (Top 10)</i>		0.632
<i>Total Features in Model</i>		86

Notes: This table reports Mean Decrease in Impurity (MDI) importance scores for the top 10 features in the first-stage Random Forest classifier. The classifier predicts whether an occupation has any potential for remote work (binary outcome). MDI measures the total reduction in node impurity (Gini coefficient) attributable to splits on each feature, averaged across all trees in the forest. Higher values indicate features that provide more discriminatory power. The dominance of communication-related features (Electronic Mail, Telephone, Face-to-Face Discussions) reflects that remote work fundamentally depends on the ability to coordinate and communicate without physical presence. The top 10 features collectively account for 63.2% of the model’s total predictive importance. The classifier was trained on 812 occupations from the BLS Occupational Requirements Survey with 86 O*NET work activity, work context, skill, and ability features.

remote work, not just worker preferences or firm policies.

Figure 10 provides a visual summary of feature importance across both stages, showing the relative contribution of O*NET dimensions (work activities, work context, skills, abilities).

B.4.3 Face Validity: Occupational Rankings

Table 8 lists occupations at the extremes of the ψ distribution. The index correctly identifies knowledge-intensive roles requiring minimal physical presence (Actuaries 95%, Software Developers 92%, Financial Analysts 88%) as having high remote work potential. Conversely, occupations requiring physical tasks or in-person service delivery (Construction Workers 5%, Healthcare Support 2%, Personal Care Workers 2%) score near zero. This face validity check confirms the index is capturing an intuitive economic concept.

Table 7: Feature Importance: Stage 2 Regressor (Intensive Margin)

Rank	O*NET Feature	MDI Importance
1	Handling and Moving Objects	0.190
2	Controlling Machines and Processes	0.078
3	Assisting and Caring for Others	0.062
4	Spend Time Standing	0.056
5	Performing General Physical Activities	0.051
6	Working with Computers	0.043
7	Spend Time Using Hands to Handle Objects	0.041
8	Inspecting Equipment, Structures, or Materials	0.038
9	Operating Vehicles or Equipment	0.035
10	Deal with Physically Aggressive People	0.032
<i>Cumulative Importance (Top 10)</i>		0.626
<i>Total Features in Model</i>		86

Notes: This table reports Mean Decrease in Impurity (MDI) importance scores for the top 10 features in the second-stage Random Forest regressor. The regressor predicts the degree of remote work potential (continuous outcome on [0,1]) for occupations classified as amenable in Stage 1. MDI measures the total reduction in mean squared error attributable to splits on each feature, averaged across all trees. The dominance of physical task requirements (Handling/Moving Objects, Controlling Machines, Physical Activities) with negative effects indicates that even among potentially teleworkable occupations, those requiring hands-on work have lower remote potential. Conversely, Working with Computers (rank 6) positively predicts higher ψ values. The top 10 features account for 62.6% of the model's predictive importance. The regressor was trained on 623 occupations classified as having remote potential (Stage 1 positives) using 86 O*NET features.

B.5 Moment Computation and Bootstrap Standard Errors

This subsection documents the computation of the 14 empirical moments used in structural estimation and the bootstrap procedure for obtaining standard errors.³

B.5.1 Target Moments


We compute 14 empirical moments per year (2019 and 2024) for structural estimation, organized into three categories:

Distributional Moments (7):

- (1) mean_logwage: Mean log real wage
- (2) var_logwage: Variance of log real wages
- (3) p90_p10_logwage: 90th-10th percentile log wage ratio
- (4) mean_alpha: Mean remote work share
- (5) var_alpha: Variance of remote work share
- (6) inperson_share: Fraction with $\alpha \leq 0.2$ (mostly in-person)
- (7) mean_psi: Mean teleworkability index

³Replication code: `src/empirical/data_moments_lean.jl` and `bootstrap/compute_bootstrap_moments_threaded.jl`

Figure 10: Feature Importance Across Classification and Regression Stages



Note: Bars show MDI importance for top 15 features in each stage. Blue = first-stage classifier, Orange = second-stage regressor.

Regression Coefficients (5):

(8) wage_psi_slope: Linear coefficient on ψ in wage regression (β_ψ)

(9) wage_psi_curvature: Quadratic coefficient on ψ^2 (β_{ψ^2})

(10) wage_alpha: Linear coefficient on α (β_α)

(11) wage_alpha_curvature: Quadratic coefficient on α^2 (β_{α^2})

(12) alpha_psi_slope: Coefficient on ψ in α regression (sorting)

External Moments (2):

(13) market_tightness: $\theta = v/u$ from FRED JOLTS and UNRATE data

(14) job_filling_rate: $q(\theta) = \theta^{-0.5}$ computed from market tightness

B.5.2 Regression Approach

Unified Regression. Rather than running separate regressions for each moment, we estimate a single comprehensive wage regression per year:

$$\log(w_i) = \beta_\psi \psi_j + \beta_{\psi^2} \psi_j^2 + \beta_\alpha \alpha_i + \beta_{\alpha^2} \alpha_i^2 + \mathbf{X}_i' \Gamma + \text{FE}_s + \text{FE}_{\text{ind}} + \epsilon_i$$

where $\mathbf{X}_i$ includes age, age², experience, experience², sex, race, and education, and we include fixed effects for state and industry. We estimate this using `FixedEffectModels.jl` with CPS survey weights (WTFINL). From this single regression, we extract coefficients β_ψ , β_{ψ^2} , β_α , and β_{α^2} as moments 8-11.

For moment 12 (sorting), we run a separate OLS regression of α on ψ with the same controls and fixed effects, extracting the coefficient on ψ .

Sample Filters. Beyond the main sample restrictions described in Appendix B.2, moment computation requires non-missing α , non-missing ψ , finite log wage, and all demographic controls present. After these filters, the moment computation samples contain 72,450 observations (2019) and 87,230 observations (2024), representing 98% retention from the main analysis sample.

B.5.3 Bootstrap Procedure

B.5.4 Bootstrap Procedure

To compute standard errors for the 14 moments, we implement a block bootstrap stratified by occupation-education cells:

Step 1: Create Bootstrap Samples (N=500):

For each bootstrap iteration $b = 1, \dots, 500$:

- Identify unique cells defined by (SOC minor group $\times$ education category)
- Sample cells with replacement (approximately 150 cells)
- Within each sampled cell, sample workers with replacement
- Construct bootstrap dataset of size equal to original sample

This block structure preserves within-cell correlation in wages and remote work patterns, which is critical for accurate standard error estimation when moments depend on occupation-level variation.

Step 2: Compute Moments on Each Bootstrap Sample:

Using parallel computation (48 cores), for each bootstrap sample b :

- Load bootstrap dataset
- Apply same sample filters as original data
- Compute all 14 moments using identical procedures
- Save moments to disk

Step 3: Aggregate Standard Errors:


Load all 500 bootstrap moment vectors. For each moment m , compute:

$$SE(m) = \sqrt{\frac{1}{500} \sum_{b=1}^{500} (m_b - \bar{m})^2}$$

where m_b is the moment value in bootstrap sample b and $\bar{m}$ is the mean across bootstraps.

Bootstrap Results. For 2019, distributional moments have SE ranging 0.002-0.015, regression coefficients have SE ranging 0.008-0.025, and market tightness has SE = 0.012. Standard errors for 2024 are similar in magnitude. These bootstrap SEs are used to construct the weighting matrix W in SMM estimation, giving more weight to precisely

Figure 11: Remote Work Patterns by Education Level (Detailed)



Note: Detailed comparison by education: (A) Remote work adoption (α), (B) Teleworkability potential (ψ), and (C) Implementation gap ($\psi - \alpha$).

estimated moments.

C Appendix Figures

Table 8: Teleworkability Index: Occupations at the Distribution Extremes

Panel A: Highest Teleworkability (ψ)	
Occupation	ψ Score
Actuaries	0.95
Software Developers	0.92
Financial Analysts	0.88
Computer and Information Research Scientists	0.87
Statisticians	0.86
Operations Research Analysts	0.84
Accountants and Auditors	0.82
Financial Managers	0.81
Mathematicians	0.80
Market Research Analysts	0.79
Panel B: Lowest Teleworkability (ψ)	
Occupation	ψ Score
Personal Care Aides	0.02
Home Health Aides	0.02
Nursing Assistants	0.03
Construction Laborers	0.05
Food Preparation Workers	0.06
Janitors and Cleaners	0.07
Automotive Service Technicians	0.08
Retail Salespersons	0.09
Truck Drivers (Heavy)	0.10
Childcare Workers	0.11

Notes: This table displays the 10 occupations with the highest and lowest teleworkability index (ψ) scores to demonstrate face validity. The index ranges from 0 (no remote work potential) to 1 (fully teleworkable). High- ψ occupations are knowledge-intensive roles requiring minimal physical presence or hands-on tasks (e.g., Actuaries, Software Developers, Financial Analysts). These positions involve primarily cognitive work that can be performed via computer and communication technology. Low- ψ occupations require physical presence for in-person service delivery, hands-on manipulation of objects, or operation of equipment/vehicles (e.g., Personal Care Aides, Construction Laborers, Food Preparation Workers). The clear economic intuition behind these rankings validates that the machine learning model captures technological feasibility rather than artifacts of the training data. Occupations are defined at the SOC 6-digit level and mapped to CPS occupation codes for analysis.