The Experimentation Value of Occupations

Piero De Dominicis*

September 28, 2026

Abstract

Early in the career, workers are uncertain about their abilities and may experiment with different jobs to find where they are most productive. In this paper I estimate a dynamic structural model of occupational choice and learning on NLSY79 panel data to quantify the experimentation value of occupations over the lifecycle. This experimentation value arises when workers choose occupations in which they do not know their ability as this may result in a good match. I show this experimentation motive acts as a compensating differential particularly early in the career, drives the mobility of young workers and declines quickly over the lifecycle. Allowing for unobserved heterogeneity, I find that the willingness to pay for learning is heterogeneous across worker types and higher for high-ability workers.

Keywords: Occupational Choice, Learning, Experimentation, Comparative Advantage, Compensating Differentials

*Department of Economics, Bocconi University, piero.dedominicis@phd.unibocconi.it

1 Introduction

Occupations differ not only in the wage they offer and the human capital they allow workers to accumulate, but also in the information they provide about workers' abilities. At the beginning of their careers, when comparative advantage is more uncertain, workers can acquire information about their abilities by experimenting across occupations. A favorable outcome makes a worker more likely to remain in an occupation, whereas an unfavorable outcome may lead her to switch and explore another alternative. Over time, however, workers accumulate human capital that is only imperfectly transferable across occupations. Switching therefore becomes increasingly costly, creating a trade-off between experimenting to discover one's comparative advantage and remaining in an occupation to accumulate specialized human capital.

In this paper, I estimate a dynamic structural model of occupational choice and learning using panel data from the US to understand the importance of this experimentation motive and to study how it changes over the lifecycle. I show that in a dynamic model of learning an experimentation motive arises even if the current-period payoff is linear in the expected wage, hence in the unknown ability. This motive stems from the dynamic sorting of workers across occupations, as the continuation value of the resulting value function is convex in mean beliefs, so positive belief realizations outweigh negative ones. Therefore, when ability uncertainty in an occupation is higher, it can generate large positive returns, more so early in the career.

To maintain tractability and ensure identification of the components of the model, I group occupations into 6 clusters using O*NET task data, and I estimate the model on US panel data from the National Longitudinal Survey of Youth 1979 (NLSY79). Methodologically, this paper relies on identification results from Arcidiacono et al. (2025) and Bunting et al. (2024). Computationally, it extends the two-stage procedure developed in De Dominicis (2025) to estimate unobserved permanent types and to allow for a type-specific prior at labor-market entry. In a first stage I estimate the wage-equation parameters and recover the distribution of unobserved types, while controlling for belief-based dynamic selection. In a second stage I estimate type-specific prior beliefs based on the observed occupational choices and the wage signals recovered in the first stage.

Using the estimated parameters and the approximated value function, I compute the experimentation value of occupations at each point in time for each worker type. Because the flow payoff is in log-wage units, the estimated experimentation value can be read as a compensating differential: the wage cut a worker would accept to work in an occupation in exchange for the information it delivers about her abilities. I find that the experimentation motive is large at labor-market entry, but quite heterogeneous across worker types. The lowest-ability type has very low belief dispersion when entering the labor market and would learn little from experimenting. Hence, workers of this type immediately sort into the occupations in which they have a known comparative advantage. These workers are quite immobile, in the sense that they switch little over the lifecycle. At labor-market entry, this willingness to pay is substantial for high-ability workers, up to 24 percent of the wage in the most informative type–occupation pairs, and 17–18 percent at the occupations these workers actually choose. For the lowest-ability type it is essentially zero over the entire career, as these workers already enter the labor market with quite narrow beliefs. The value of experimenting is highest at the beginning of the career, fading to essentially zero after roughly fifteen years in the labor market, and it leads workers to enter occupations they later leave once they learn they are not a good match.

These findings contribute to the literature studying the causes and consequences of early-career mobility. In particular, I propose an experimentation motive in which workers try out occupations when entering the labor market and then switch when the realized productivity component is not good enough. Models of learning and job mobility were proposed initially by Jovanovic (1979) and Miller (1984), and more recently by Antonovics and Golan (2012), Papageorgiou (2014), Pastorino (2024), and De Dominicis (2025). Using a structural model of learning about occupation-specific ability, I show that an experimentation motive arises even if workers behave as risk-neutral in the current-period expected utility: it comes purely from a dynamic sorting motive. Choosing an uncertain occupation comes with an option value: if the outcome realization is bad the worker can switch occupation, while if the outcome is good then the worker remains in that occupation. Additionally, I allow for permanent unobserved heterogeneity in types and show that these different types value experimentation differently, hence they have very diverse realized occupational mobility patterns.

2 Data and Occupation Clusters

2.1 Sample construction

To estimate the model I use panel data from the National Longitudinal Survey of Youth 1979 (NLSY79). To construct the sample, I follow the approach used in Guvenen et al. (2020). The resulting sample selection defines a yearly panel of male workers with valid hourly wages, occupational tenure, three-digit Census occupation codes, and pre-market ability test scores from the 1980 administration of the Armed Services Vocational Aptitude Battery. Importantly, these test scores represent measurements of abilities that are independent of labor-market outcomes. The main measure of wage used in the paper is obtained by deflating nominal hourly wages with the CPI deflator. I restrict attention to individuals whose age at the first interview is between 16 and 24.

The estimation sample contains a total of $N = 1{,}871$ male workers and $NT = 43{,}342$ person-year observations, covering the years 1978–2010 and ages 16–60. The average worker in the sample is observed for roughly twenty-three years out of the available thirty-three, holds about four distinct three-digit Census occupations over the lifecycle, and earns a mean log hourly wage of 7.28. Workers in the sample are 11.3% Black and 6.7% Hispanic, and on average have 13.6 years of schooling. Summary statistics for all variables entering the wage and measurement blocks are reported in Table 2 in Appendix A.

2.2 Clustering Occupations

Mainly for identification and estimation purposes, but also to improve the interpretability of the estimation results, I group three-digit occupations into $J = 6$ clusters. Clusters are defined using a k-means algorithm over the 120 standardized O*NET descriptors, computed at the occupation level and weighted by worker-year employment mass. Details of the construction, robustness, and the full list of top occupations by cluster are in Appendix A.2. Table 1 reports the resulting clusters, together with the share of worker-year observations, the average log wage and the average tenure in each cluster. Clusters are labeled using the job titles of their most common occupations. Intuitively, these clusters capture not only horizontal and vertical differentiation of occupations, but also differences in task content (mainly cognitive, manual and interpersonal) and in their relevance across economic sectors. For example, cluster $j = 0$ mainly includes manual jobs at the bottom of the occupational ladder in the manufacturing sector, while cluster $j = 2$ groups transport and protective-service occupations, such as truck drivers and guards. Cluster $j = 3$ captures high-skill technical and engineering roles, and cluster $j = 5$ captures management and professional roles, such as managers, accountants and lawyers.

j	Label	Share	Log wage	Tenure	Representative occupations
0	Manual, laborer	0.259	7.08	8.3	machine operators, carpenters, laborers
1	Skilled repair, craft	0.088	7.15	6.7	auto mechanics, electricians, machinists
2	Transport, protective	0.075	7.06	6.3	truck drivers, guards, pilots
3	Technical, engineering	0.151	7.55	8.1	software developers, engineers, foremen
4	Sales, clerical, services	0.179	7.14	7.0	salespersons, cooks, police, clerks
5	Managerial, professional	0.249	7.55	8.8	managers, accountants, lawyers

Table 1: "Share" is the worker-year employment share, "Log wage" the mean log hourly wage, "Tenure" the mean occupation-cluster tenure in years. Source: NLSY79 + O*NET. See Appendix A.2 for more details on the clustering procedure.

3 Structural Model

The economy is populated by a continuum of workers. Time is discrete; in each period t a worker i chooses an occupation $j \in \mathcal{O} = \{0, 1, \ldots, J-1\}$. The model is closely related to De Dominicis (2025), but differs in a few key aspects. Following Arcidiacono et al. (2025), I introduce latent types that shift both the overall level of wages and the variance of the unobserved abilities across occupations.

3.1 Types, abilities and the wage equation

Worker i is characterized by a discrete type $r_i \in \{1, \ldots, R\}$, drawn at labor-market entry with probabilities $\pi_r = \Pr(r_i = r)$, and by a vector of continuous occupation-specific abilities $\boldsymbol{\xi}_i = (\xi_{i1}, \ldots, \xi_{iJ})' \in \mathbb{R}^J$. Conditional on the type, abilities are drawn from a type-specific multivariate normal distribution,

$$\boldsymbol{\xi}_i \mid r_i \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_{r_i}), \tag{1}$$

with $\mathbf{\Sigma}_r$ a full, unrestricted $J \times J$ covariance matrix for each discrete type r. The log wage of worker i employed in occupation cluster j at time t is

$$w_{ijt} = \alpha_j + \gamma_{j,r_i} + \xi_{ij} + F_j(\tau_{ijt}) + H_j(E_{it}) + \varepsilon_{ijt}, \qquad \varepsilon_{ijt} \sim \mathcal{N}(0, \varsigma^2_{\varepsilon j}), \tag{2}$$

where α_j is an occupation fixed effect, γ_{j,r_i} is an occupation-by-type wage premium, ξ_{ij} is the worker's latent ability in that occupation, and ε_{ijt} is an idiosyncratic productivity shock, independently and identically distributed over workers and time, with occupation-specific variance. Returns to experience operate through two occupation-specific functions: $F_j(\cdot)$ determines the returns to occupation-specific tenure τ_{ijt}, and $H_j(\cdot)$ determines the returns to total labor-market experience E_{it},

$$E_{it} \;=\; \sum_{j \in \mathcal{O}} \tau_{ijt}, \tag{3}$$

with $F_j(0) = H_j(0) = 0$ for every j. Occupational tenure and labor-market experience evolve deterministically: collecting the J tenures in $\boldsymbol{\tau}_{it} = (\tau_{i1t}, \ldots, \tau_{iJt})'$,

$$\boldsymbol{\tau}_{i,t+1} = \boldsymbol{\tau}_{it} + \boldsymbol{e}_{d_{it}}, \tag{4}$$

where $d_{it} \in \mathcal{O}$ is the chosen occupation and $\boldsymbol{e}_j$ is the j-th unit vector: a year of work adds one year of tenure in the chosen occupation and nothing in other occupations. In this model, switching occupation leads to an implicit cost due to the loss of occupation-specific returns, but no loss of labor-market experience. Therefore, E_{it} represents the fully transferable part of human capital, while τ_{ijt} is occupation-specific.

3.2 Learning and information structure

The worker knows her type r_i at entry but not her ability vector $\boldsymbol{\xi}_i$, which she learns from realized wages. All components of the wage equation (2) other than $\xi_{ij} + \varepsilon_{ijt}$ are known to the worker, so after working in occupation j at time t she observes the signal

$$s_{ijt} \;\equiv\; w_{ijt} - \alpha_j - \gamma_{j,r_i} - F_j(\tau_{ijt}) - H_j(E_{it}) \;=\; \xi_{ij} + \varepsilon_{ijt} \;=\; \boldsymbol{e}_j'\boldsymbol{\xi}_i + \varepsilon_{ijt}, \tag{5}$$

that is, a noisy measurement of ability in the current occupation. At labor-market entry, a worker of type r holds the type-specific subjective prior beliefs

$$\tilde{\boldsymbol{\xi}}_{i0} = \boldsymbol{b}_{r_i}, \qquad \tilde{\mathbf{\Sigma}}_{i0} = \mathbf{\Sigma}^s_{r_i} \equiv e^{h_{r_i}}\, \mathbf{\Sigma}_{r_i}, \tag{6}$$

where $\boldsymbol{b}_r \in \mathbb{R}^J$ is a vector of subjective prior means, normalized to mean zero across occupations within each type ($\mathbf{1}'\boldsymbol{b}_r = 0$). This normalization restricts the average prior belief across occupations to equal the population mean of ability, so that $\boldsymbol{b}_r$ captures only perceived comparative advantage: different types can enter the labor market with different beliefs about their relative abilities across occupations. For the prior beliefs covariance matrix, I adopt a parsimonious parametrization with one free scale per type h_r, so the subjective correlation structure across occupations is inherited from the objective covariance $\mathbf{\Sigma}_r$. This scale parameter can increase or reduce dispersion of prior beliefs at entry. Rational expectations is nested at $\boldsymbol{b}_r = \mathbf{0}$ and $h_r = 0$.

A distinctive feature of the model with occupation-specific learning structure is that the subjective posterior covariance is a deterministic function of the occupational tenure vector. Each year in an occupation adds the occupation's perceived signal precision to the corresponding diagonal entry of the posterior information matrix, so that after any tenure history,

$$\tilde{\boldsymbol{\Sigma}}_{it} \;=\; V^s_{r_i}(\boldsymbol{\tau}_{it}) \;\equiv\; \Big(\big(\boldsymbol{\Sigma}^s_{r_i}\big)^{-1} + \sum_{j \in \mathcal{O}} \frac{\tau_{ijt}}{\varsigma^2_{\varepsilon j}}\, \boldsymbol{e}_j \boldsymbol{e}_j' \Big)^{-1}, \tag{7}$$

regardless of the order in which the signals arrived. The covariance therefore never needs to be carried as a separate state variable: the tenure vector, already a state through the wage equation, encodes the entire uncertainty stock. This drastically reduces the state variables of the dynamic program when the number of occupations grows. Upon observing the wage signal in the chosen occupation, the belief mean is updated by the subjective Kalman recursion

$$\tilde{\boldsymbol{\xi}}_{i,t+1} = \tilde{\boldsymbol{\xi}}_{it} + \boldsymbol{K}_{ijt}\big(s_{ijt} - \boldsymbol{e}_j' \tilde{\boldsymbol{\xi}}_{it}\big), \qquad \boldsymbol{K}_{ijt} = \frac{V^s_{r_i}(\boldsymbol{\tau}_{it})\, \boldsymbol{e}_j}{\boldsymbol{e}_j' V^s_{r_i}(\boldsymbol{\tau}_{it})\, \boldsymbol{e}_j + \varsigma^2_{\varepsilon j}}, \tag{8}$$

where the realized signal s_{ijt} is defined in (5). Two forces shape the learning dynamics in this model. First, the prior correlation structure drives belief spillovers across occupations: a signal in one occupation moves beliefs about every other occupation in proportion to their covariance. Because the variance-covariance matrix is type specific, so is the cross-occupation information spillover: a type whose abilities are strongly correlated, either positively or negatively, updates its beliefs about unobserved abilities in many occupations, while a type with nearly independent abilities must try different occupations to learn. Second, the wage noise: noisier occupations have lower signal-to-noise ratios and reveal less information about workers' abilities.

3.3 Occupational choice

The worker state at the beginning of period t is $x_{it} \equiv \big(r_i, \tilde{\boldsymbol{\xi}}_{it}, \boldsymbol{\tau}_{it}\big)$ and includes her type r_i, the current subjective belief about her ability vector $\tilde{\boldsymbol{\xi}}_{it}$, and the vector of her current tenures $\boldsymbol{\tau}_{it}$. In each period t worker i chooses an occupation $j \in \mathcal{O}$ before observing the wage realization in that occupation. I assume log utility, $u(W) = \log W$, so that the current-period expected utility is:

$$\mathbb{E}\big[\log W_{ijt} \mid x_{it}\big] = \alpha_j + \gamma_{j,r_i} + \tilde{\xi}_{ijt} + F_j(\tau_{ijt}) + H_j(E_{it}), \tag{9}$$

where W_{ijt} is the wage paid in levels. At the end of the period the worker observes the wage, tenure accumulates by (4), and beliefs update by (8). Figure 1 below summarizes the within-period timing of the worker's choice and the resulting belief updating.

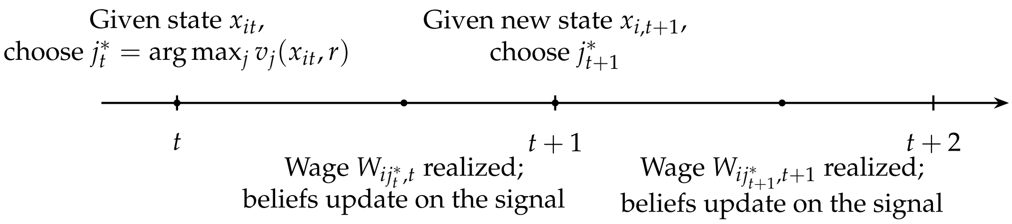

Figure 1: Timing of choices and wage realizations.

Dropping the worker index, the worker's dynamic program is

$$V^r(\tilde{\boldsymbol{\xi}}_t, \boldsymbol{\tau}_t) = \max_{j \in \mathcal{O}} \Big\{ \mathbb{E}\big[\log W_{jt} \mid \tilde{\boldsymbol{\xi}}_t, \boldsymbol{\tau}_t\big] + \beta\, \mathbb{E}_{s_{jt}}\big[V^r(\tilde{\boldsymbol{\xi}}_{t+1}, \boldsymbol{\tau}_{t+1})\big] \Big\},$$

$$\begin{aligned} \text{subject to:} \quad \boldsymbol{\tau}_{t+1} &= \boldsymbol{\tau}_t + \boldsymbol{e}_j, \\ s_{jt} &= \boldsymbol{e}_j' \boldsymbol{\xi} + \varepsilon_{jt}, \\ \tilde{\boldsymbol{\xi}}_{t+1} &= \tilde{\boldsymbol{\xi}}_t + \boldsymbol{K}_{jt}(s_{jt} - \boldsymbol{e}_j' \tilde{\boldsymbol{\xi}}_t), \\ \tilde{\boldsymbol{\Sigma}}_{t+1} &= V_r^s(\boldsymbol{\tau}_{t+1}), \\ \text{given} \quad (\tilde{\boldsymbol{\xi}}_0, \tilde{\boldsymbol{\Sigma}}_0, \boldsymbol{\tau}_0) &= (\boldsymbol{b}_r, \boldsymbol{\Sigma}_r^s, \boldsymbol{0}) \text{ for an entrant of type } r, \end{aligned} \tag{10}$$

where $\beta \in (0, 1)$ is the discount factor. The flow and the continuation value in the Bellman equation are calculated over the worker's subjective predictive distribution of the wage signal,

$$s_{jt} \mid x_t \sim \mathcal{N}\big(\boldsymbol{e}_j' \tilde{\boldsymbol{\xi}}_t,\, \boldsymbol{e}_j' V_r^s(\boldsymbol{\tau}_t) \boldsymbol{e}_j + \varsigma_{\varepsilon j}^2\big)$$

which drives the Bayesian update of beliefs. The optimal policy is the following choice rule:

$$d_t = \arg\max_{j \in \mathcal{O}} \; v_j(x_t, r), \tag{11}$$

where $v_j(x, r)$, the term in braces in (10), is the conditional value function of choosing occupation j given the state x and the type r.

4 Identification and Estimation

The parameters of the structural model are organized in three blocks,

$$\begin{aligned} \boldsymbol{\Theta}_1 &= \Big\{ (\boldsymbol{\mu}_r)_{r=1}^R,\ \boldsymbol{\Sigma}_\nu \Big\}, \\ \boldsymbol{\Theta}_2 &= \Big\{ (\alpha_j, \gamma_{j,r}, F_j, H_j, \varsigma_{\varepsilon j}^2)_{j \in \mathcal{O}},\ (\boldsymbol{\Sigma}_r)_{r=1}^R,\ (\pi_r)_{r=1}^R \Big\}, \\ \boldsymbol{\Theta}_3 &= \Big\{ (\boldsymbol{b}_r,\ h_r)_{r=1}^R \Big\}, \end{aligned}$$

collecting respectively the parameters from the measurement block, the wage block and the subjective-belief block. Identification proceeds in the same two steps as Arcidiacono et al. (2025). First, I identify the latent-type parameters $\boldsymbol{\Theta}_1$ and the wage-equation and ability-distribution parameters $\boldsymbol{\Theta}_2$ using the independent measurements and the observed wages;

then, conditional on those, I identify the parameters driving learning using data on realized choices.

4.1 Identification of the discrete types with measurement system

The ASVAB scores can be used to identify types as they provide selection-free measurements that are independent of wages and choices. The measurement block (12) defined below is a finite Gaussian mixture with R components, diagonal within-component covariance, and seven conditionally independent indicators. By Theorem 8 of Allman et al. (2009), three conditionally independent indicators are enough for identification of a finite mixture with distinct component means; with seven indicators the type means, the measurement noise variances, and the mixture probabilities are identified up to label permutation. Therefore, I order types by their mean test score. The correlation structure of the seven tests clusters the population into R discrete types without using information about labor-market outcomes; the wage observations then reveal the premium these types earn in each occupation. Intuitively, measurements identify types, while wages identify premiums.

More specifically, the discrete types are identified from a measurement block of $P = 7$ age-residualized, standardized ASVAB subtests $\boldsymbol{y}_i = (y_{i1}, \ldots, y_{iP})'$, taken before labor-market entry: Word Knowledge, Paragraph Comprehension, Arithmetic Reasoning, Mathematics Knowledge, Mechanical Comprehension, Auto & Shop Information, and Electronics Information. Conditional on the type, the scores are jointly normal with a free mean vector and diagonal noise,

$$\boldsymbol{y}_i \mid r_i \sim \mathcal{N}(\boldsymbol{\mu}_{r_i}, \boldsymbol{\Sigma}_\nu), \qquad \boldsymbol{\Sigma}_\nu = \operatorname{diag}(\sigma^2_{\nu 1}, \ldots, \sigma^2_{\nu P}). \tag{12}$$

Because these tests are taken before entering the labor market, $\boldsymbol{y}_i$ is independent of wages, occupational choices, and $\boldsymbol{\xi}_i$ conditional on the type r_i.

4.2 Identification of the wage block under learning

Conditional on types, identification of (2) in the presence of endogenous occupational choice follows the pure-learning selection argument of Bunting et al. (2024), in the form used by Arcidiacono et al. (2025) and De Dominicis (2025). The model is a pure-learning environment: workers sort on beliefs, and beliefs are a deterministic function of the same wage and occupational histories the econometrician observes, conditional on the type. Letting $\mathcal{H}_{it} \equiv (r_i, D_i^{t-1}, W_i^{t-1}, \boldsymbol{\tau}_i^t)$ denote the worker's history, the model implies the following conditional-independence restriction

$$D_{it} \perp\!\!\!\perp \boldsymbol{\xi}_i \mid \mathcal{H}_{it},$$

so that, conditional on history, selection is on observables and the type-conditional potential-wage path is recovered under any supported occupational sequence. I now provide the intuition for identification; for a more formal treatment the reader can refer to Bunting et al. (2024) and De Dominicis (2025).

The type-conditional mean of the potential wage path identifies the sum of the occupation fixed effect and the type premium, and normalizing one reference type in each occupation[1] separates the two. Within-type covariances of weighted wage residuals identify the ability covariances: repeated observations in the same occupation separate the permanent ability component from the transitory noise, identifying the ability variances type by type, while workers observed in two clusters identify the corresponding ability covariances from the cross-occupation covariance of their residuals. Type-specificity of the covariances is identified because the probability of being of a certain type is pinned down by the selection-free measurement block, so within-type second moments can also be recovered. The tenure and experience functions are identified from within-worker within-occupation wage growth: general experience grows every year while occupation-cluster tenure grows only in years worked in the cluster, so continuing spells identify the sum of the two components while movers separate them.

4.3 Identification of the dynamic belief parameters

The measurement system and realized wages identify the distribution of unobserved types and the wage equation parameters. Conditional on these Stage-1 estimates, the remaining parameters $((\boldsymbol{b}_r, h_r)_{r=1}^R)$ determine workers' type-specific subjective prior beliefs,

$$\boldsymbol{\xi}_i \mid r \sim_s \mathcal{N}\left(\boldsymbol{b}_r, e^{h_r}\widehat{\boldsymbol{\Sigma}}_r\right),$$

where $\widehat{\boldsymbol{\Sigma}}_r$ is the type-specific covariance of permanent occupational abilities identified from the wage panel. The prior mean vector $\boldsymbol{b}_r$ is normalized so that its average across occupations is zero and therefore captures the perceived comparative advantage of types across occupations, rather than absolute advantage. The mean prior beliefs about unobserved abilities $\boldsymbol{b}_r$ are identified primarily from occupational choices at labor-market entry: if workers of a given type enter a particular occupation more frequently than the expected wages implied by the Stage-1 estimates would predict, the model attributes this excess sorting to a higher prior mean belief about productivity in that occupation. On the other hand, the parameter h_r governs the scale of perceived initial uncertainty. Because both $\widehat{\boldsymbol{\Sigma}}_r$ and the occupation-specific transitory wage-noise variances are already identified from wages, h_r is identified by the observed mobility patterns over the lifecycle. The initial prior uncertainty determines both (i) the updating of mean beliefs over the career and (ii) the experimentation value of different occupations at the beginning of the career. For example, greater perceived initial uncertainty implies stronger posterior revisions following wage realizations and, consequently, more occupational mobility over the lifecycle. At the same time, more uncertainty also increases the experimentation value: if workers enter an occupation even though its expected wage is relatively low, they must value the potential positive productivity realizations in that occupation.

4.4 Two-stage estimation

In De Dominicis (2025) I show how to write down the joint likelihood of observed wages and choices and how a useful factorization allows a similar version of the model presented in this paper to be consistently estimated in two steps. However, in the model presented in Section 3,

[1] $\gamma_{j,1} = 0$ for type 1 workers in each occupation j.

there are two important deviations that require modifications of the estimation strategy. First, and most importantly, the unobserved types now have to be identified and recovered, which I do using selection-free measurements. Second, the occupational choice is deterministic, in the sense that, conditional on the state variables, there is no random shock that makes the occupational choice probabilistic. This makes estimation much harder, as the conditional choice likelihood, given workers' states, would be degenerate. In fact, beliefs are a deterministic function of the observed wage history and the type, so conditional on the data the model predicts a specific choice path. In the model, workers face genuine uncertainty about their future wage draws, but none of the randomness that drives their choices is hidden from the econometrician: since the sequence of wages is also observed, the exact sequence of beliefs can be computed. Hence, the standard conditional choice probabilities (CCP) used for estimation in discrete choice models are not available.

I therefore estimate the model with a modified two-step procedure. The first stage maximizes the likelihood of wages and test scores conditional on the observed occupation paths using an expectation-maximization algorithm, and the second stage estimates the type-specific prior distributions by simulated method of moments. Conditioning on the occupation paths in the first stage is valid because of the information structure of the model. Workers never observe their abilities: they choose occupations based on their beliefs, and beliefs are a function only of the observed wage history and the type. Conditional on the wage history and the type the occupation path carries no additional information about the abilities or the wage shocks. Moreover, given the timing of the model, the choice in period t is made before the wage shock in period t is realized. Selection into occupations is based on observables, and the first-stage likelihood can be used to estimate $(\boldsymbol{\Theta}_1, \boldsymbol{\Theta}_2)$. The Stage-1 observed-data log-likelihood is

$$\ell_1(\boldsymbol{\Theta}_1, \boldsymbol{\Theta}_2) = \sum_{i=1}^{N} \log\left[\sum_{r=1}^{R} \pi_r \, L_i^{w,\text{marg}}(\boldsymbol{w}_i \mid r; \boldsymbol{\Theta}_2) \, L_i^{m}(\boldsymbol{y}_i \mid r; \boldsymbol{\Theta}_1)\right], \tag{13}$$

with $L_i^{w,\text{marg}}$ the Gaussian wage marginal obtained by integrating $\boldsymbol{\xi}_i \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_r)$ out of L_i^{w} in closed form.

4.4.1 Stage 1: EM algorithm

In the estimation, I parametrize the functions F_j and H_j with stepwise returns[2] with yearly bins for the first nine years, two-year bins from ten years up to the cap, and one pooled open-ended top bin at the cap; both profiles are normalized to zero at zero seniority, $F_j(0) = H_j(0) = 0$. I prefer this specification to quadratic returns in order not to impose decreasing late-career returns that the data cannot measure properly, as the sample does not include many late-career observations. I compare the stepwise functions with smooth quadratic returns in Appendix Figure 9. The E-step computes, for every worker-type pair, the Gaussian posterior of the ability vector and the posterior type weight. Stacking worker i's wage panel and letting $\Lambda_i \in \mathbb{R}^{T_i \times J}$

[2] Capped at the 90th percentile of the respective distributions (18 years of occupational tenure and 24 years of experience), to avoid basing late-career returns on few observations and generating unrealistic mobility patterns.

collect the unit-loading rows $\boldsymbol{e}'_{d_{it}}$ and $\Omega_i = \operatorname{diag}(\varsigma^2_{\varepsilon,d_{it}})$,

$$V_{ir} = \left(\boldsymbol{\Sigma}_r^{-1} + \Lambda_i'\Omega_i^{-1}\Lambda_i\right)^{-1}, \qquad \boldsymbol{m}_{ir} = V_{ir}\,\Lambda_i'\Omega_i^{-1}\left(\boldsymbol{w}_i - \boldsymbol{\mu}_i(r)\right), \tag{14}$$

where $\boldsymbol{\mu}_i(r)$ is the deterministic wage mean of type r. The posterior covariance V_{ir} is type-specific because the ability distribution $\boldsymbol{\Sigma}_r$ is. The updating formula for the posterior type weight follows Bayes' rule for discrete types

$$\omega_{ir} = \frac{\pi_r\, L_i^{w,\text{marg}}(\boldsymbol{w}_i \mid r)\, L_i^m(\boldsymbol{y}_i \mid r)}{\sum_{s=1}^R \pi_s\, L_i^{w,\text{marg}}(\boldsymbol{w}_i \mid s)\, L_i^m(\boldsymbol{y}_i \mid s)}, \tag{15}$$

evaluated in logs for numerical stability.

The M-step takes the E-step posteriors over abilities and types as given and computes the updated parameter vector that maximizes the expected complete-data log-likelihood. All the updates can be computed in closed form as the maximization essentially amounts to weighted least squares and weighted averages. The explicit updating formulas are in Appendix D. The loop runs until the log-likelihood increment falls below 10^{-7}, from multiple random starting points to avoid getting stuck in local optima, selecting the run with the highest converged likelihood.

4.4.2 Stage 2: SMM algorithm

In the second stage I estimate the belief parameters $\boldsymbol{\Theta}_3 = (\boldsymbol{b}_r, h_r)_{r=1}^R$ by simulated method of moments, given $(\widehat{\boldsymbol{\Theta}}_1, \widehat{\boldsymbol{\Theta}}_2)$ and the Stage-1 type posteriors. Let $\widehat{\boldsymbol{M}}$ denote a vector of moments computed from the panel and $\boldsymbol{M}(\boldsymbol{\Theta}_3)$ its model-simulated counterpart. The moments are occupation-switching hazards and occupational employment shares over the lifecycle, weighted by the Stage-1 type posteriors. The estimator is then

$$\widehat{\boldsymbol{\Theta}}_3 = \arg\min_{\boldsymbol{\Theta}_3} \left[\widehat{\boldsymbol{M}} - \boldsymbol{M}(\boldsymbol{\Theta}_3)\right]' \widehat{\boldsymbol{W}} \left[\widehat{\boldsymbol{M}} - \boldsymbol{M}(\boldsymbol{\Theta}_3)\right], \tag{16}$$

with $\widehat{\boldsymbol{W}}$ a positive-definite weighting matrix.

Each criterion evaluation nests the solution of the worker dynamic problem. Recall the state $x \equiv (\tilde{\boldsymbol{\xi}}, \boldsymbol{\tau})$. For each type r, the inner step solves for the value function $V^r(\,\cdot\,; \widehat{\boldsymbol{\Theta}}_2, \boldsymbol{\Theta}_3)$ as the fixed point of the Bellman operator implied by (10):

$$V^r(x) = \max_{j \in \mathcal{O}} \Big\{ \mathbb{E}\left[\log W_{jt} \mid x\right] + \beta\, \mathbb{E}\left[\, V^r(x') \mid x, d = j\right] \Big\}, \tag{17}$$

where the expectation over the continuation value is taken over the subjective predictive distribution of the wage signal under choice j. That integration is computed by one-dimensional Gauss–Hermite quadrature along the Kalman updating ray generated by that signal (19). Because the choice is discrete, the value function is kinked at belief indifference points, and the solution and approximation methods differ from the case when taste shocks smooth the integrated value function[3]. While De Dominicis (2025) uses value function approximation with

[3]For example, in the presence of extreme-value type-I taste shocks, the max operator is smoothed into the log-sum-exp operator: the larger the dispersion of the taste shocks, the smoother the integrated

Chebyshev polynomials, here I resort to a more general approximation method that better handles kinks. In practice, I use the adaptive sparse-grid method of Brumm and Scheidegger (2017). The value function is approximated with piecewise-linear basis functions on a grid that adaptively adds points where the current approximation is least accurate, so that the resolution concentrates exactly around the points where the kinks are more problematic. Rather than solving on a discretized state space, I solve for type-specific value functions defined over the continuous beliefs and tenure vector, and evaluate them by interpolation at any state reached in the simulation.

Given the solved fixed points I simulate the model and generate a synthetic panel, which I use to compute occupational mobility and type-specific employment shares over the lifecycle. The procedure alternates between (i) solving the fixed point in (17) for each V^r at the current $\mathbf{\Theta}_3$ and (ii) updating $\mathbf{\Theta}_3$ from the criterion (16), until convergence. Computational details are in Appendix E.

4.4.3 Externally set parameters

The discount factor is set to $\beta = 0.9$. The number of types is set to $R = 4$; the number of occupation clusters $J = 6$ is fixed by the task-space clustering of Section 2.2. All remaining parameters are estimated.

5 Estimation Results

5.1 First Stage Estimation Results

I now turn to the estimated parameters. While all the model parameters are reported in Appendix Table 4, in this section I discuss the wage levels and wage premia for the different types, the returns to experience and the type-specific unobserved ability distribution.

Given the normalization $\gamma_{j,1} = 0$ for all occupations, α_j can be interpreted as determining the wage level of type 1 in each occupation, and $\gamma_{j,r}$ for $r \in \{2, 3, 4\}$ as the wage premium of the other types with respect to type 1. Therefore, the overall comparative advantage and wage level in different occupations at labor-market entry are defined, for each type, by the sum of the occupation fixed effect and the type wage premium, $\alpha_j + \gamma_{j,r}$. I plot the sum of these two components in the left panel of Figure 2. The type schedules show that the labeling of the types, which was defined on average test scores, is also respected in terms of average wage levels across occupations; the gap between the top and the bottom type ranges from about 0.31 log points in the transport and protective services cluster to about 0.64 in the managerial and professional cluster. This is not an artifact of the imposed labeling, whose basis is plotted in the right panel of Figure 2: individuals with higher average test scores have in fact higher average wage levels in the labor market, independently of the occupation they sort into[4].

In Figure 3 I report the estimated returns to occupational tenure and total labor-market experience for different occupations. Two main messages emerge from a broad view of these

value function.

[4]Presumably, test scores are good predictors of the ability level in the labor market.

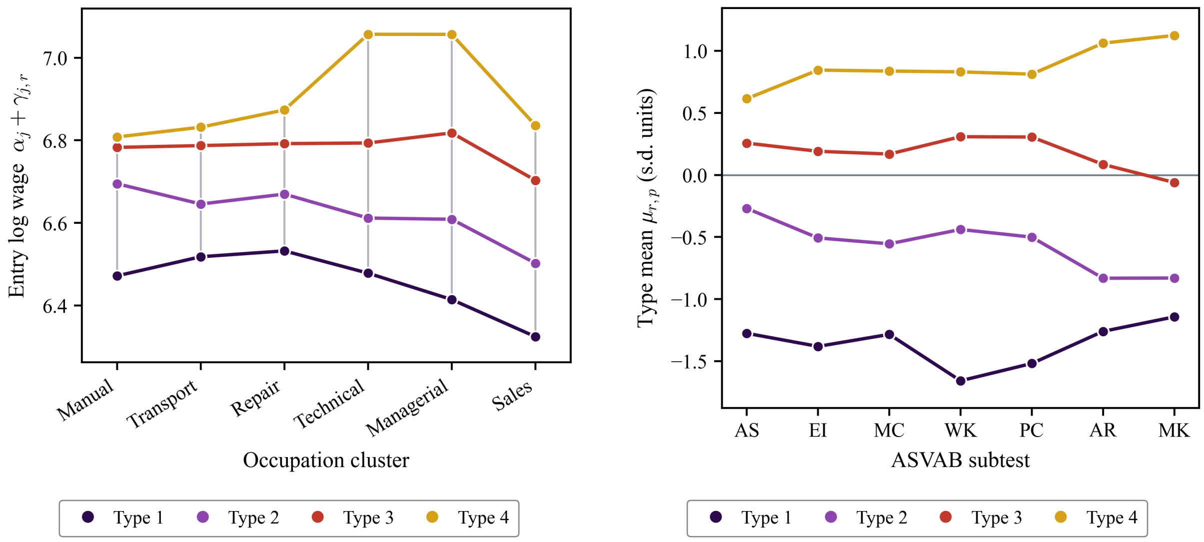

Figure 2: *Type entry wage schedules and type test-score profiles.* Left: entry log wage by occupation cluster for each type, the sum of the cluster intercept and the type premium, which is invariant to the type-reference normalization (the premium alone is identified only relative to the reference type). Right: average test scores on the seven (residualized by age and standardized) ASVAB subtests for each type; subtests are grouped into the manual-technical block (AS, EI, MC) and the cognitive block (WK, PC, AR, MK), mirroring the occupation ordering of the left panel. Subtest abbreviations, in axis order: AS = Auto & Shop Information, EI = Electronics Information, MC = Mechanical Comprehension, WK = Word Knowledge, PC = Paragraph Comprehension, AR = Arithmetic Reasoning, MK = Mathematics Knowledge. Types are ordered by ability (Type 1 lowest). Point estimates in Appendix Table 4.

estimates. First, general labor-market experience has steeper and higher returns compared to occupation-specific human capital: at the caps, the general-experience functions reach 0.38–0.65 log points against 0.13–0.42 for the tenure premia. Second, the dispersion of returns across occupations increases with both occupational tenure and labor-market experience. In particular, the technical and engineering cluster and the managerial and professional cluster offer the largest returns to occupation-specific human capital (0.34 and 0.42 log points at the tenure cap), as well as to total labor-market experience, while the transport and protective services cluster has the lowest returns to experience and, together with skilled repair and craft, the lowest returns to occupational tenure. Sales, clerical and personal services have instead medium returns to tenure and quite large returns to experience. Moving after having accumulated experience in other occupations still allows workers to earn a premium, suggesting that human capital in these occupations is quite transferable.

From the estimation of the type-specific unobserved ability distribution, a few things are worth emphasizing. First, abilities across occupations are quite correlated for all types, although this correlation shows some heterogeneity. The average off-diagonal correlations range from $+0.39$ to $+0.52$ across types. The strongest pairwise correlations reach $+0.85$ (mechanics and transport for the lowest type) and $+0.84$ (manual trades and mechanics for the second type). The main exception is the managerial and professional cluster: the sign of management's correlations differs across types: they are positive for the second type (up to $+0.80$), but negative (-0.19) with manual trades and transport for the lowest type. Ability dispersion is largest for the highest-ability type and, within types, for the managerial and professional cluster: total ability variance is 0.87, 0.59, 0.88 and 1.25 across the ability-ordered types, and the management cluster has the largest own standard deviation in every type except Type 3, where it is

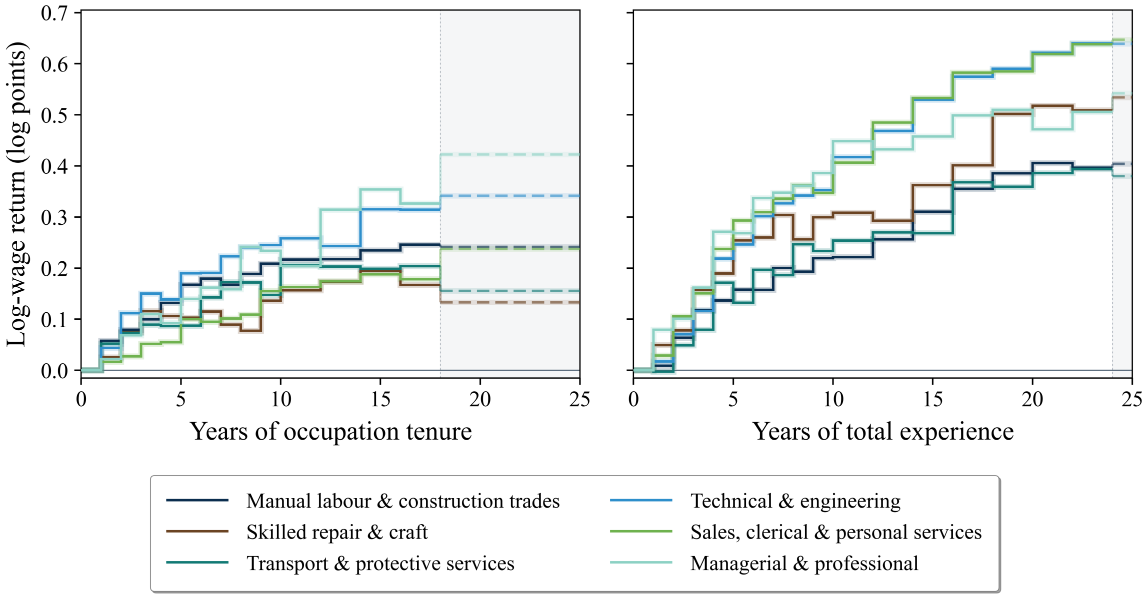

Figure 3: *Estimated returns to occupation-cluster tenure and general experience*, by cluster. The left panel plots the estimated returns to occupational tenure $F_j(\tau)$ and the right panel the general-experience functions $H_j(E)$, on shared axes from 0 to 25 years. Dashed segments beyond the caps (18 years of tenure, 24 years of experience) show how returns stay constant after a certain level of experience. Selected step coefficients in Appendix Table 4.

marginally below the technical cluster (0.38–0.58 across types), peaking at 0.58 for the highest-ability type.

5.2 Second Stage Estimation Results

In the second stage I estimate the parameters of the type-specific subjective prior distributions. The estimated values of h_r are -2.82, -1.33, -1.18 and -1.47 for Types 1–4 respectively, implying subjective prior variances of 6%, 26%, 31% and 23% of the estimated ability distribution. These parameters imply that workers have lower belief dispersion than the true ability distribution. In particular, low-type workers have very low perceived belief dispersion when entering the labor market. The estimated prior means $\boldsymbol{b}_r$, the perceived comparative advantage at entry, are reported in Appendix Table 7. Perceived comparative advantage predominantly makes workers optimistic about their abilities in more manual occupations and rather pessimistic in technical occupations. To improve interpretation of these results, I report in Figure 4 the estimated ability distribution for two specific occupations, manual trades and managerial and professional, together with the estimated type-specific prior beliefs. Because subjective beliefs are parametrized as a scaled version of the true multidimensional ability distribution, they have the same shape as the estimated ability distributions. Besides the scale, the location of the distribution also differs, as the model allows for type-specific perceived comparative advantage. These estimated perceived priors have important implications for occupational mobility over the career.

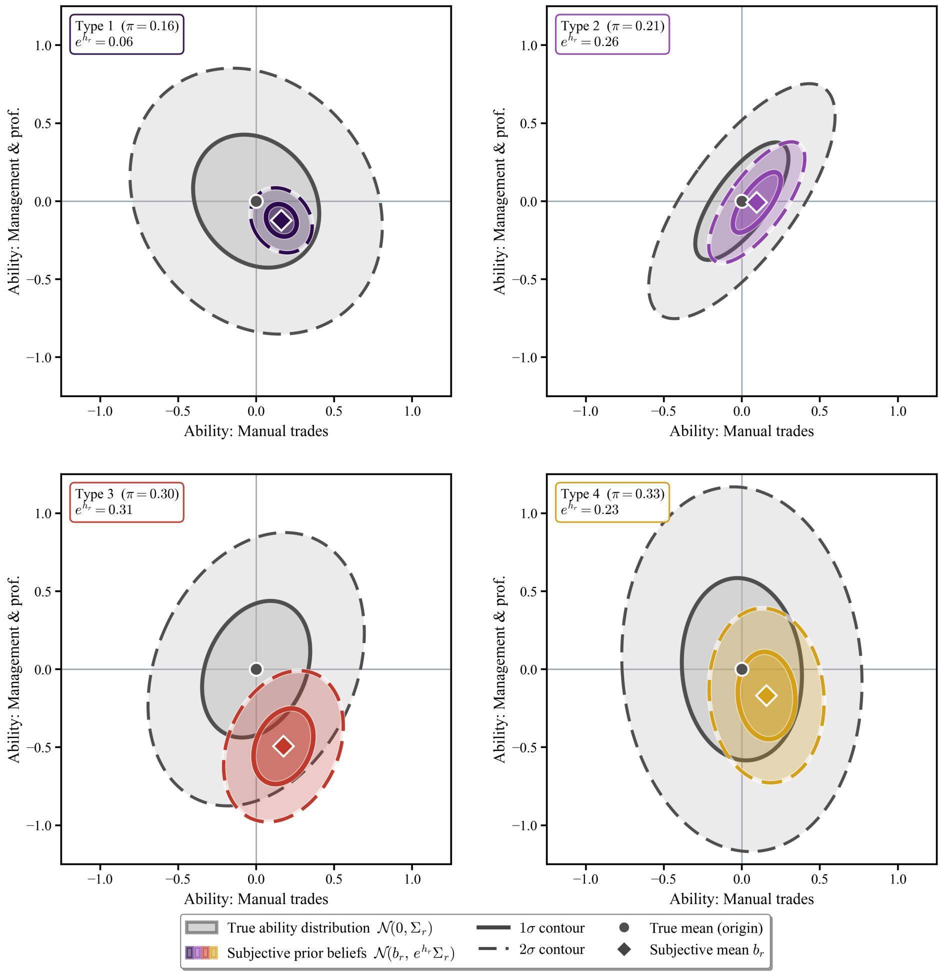

Figure 4: Estimated subjective prior beliefs and the objective ability distribution, by type. Each panel plots, for one type, the 1σ (solid) and 2σ (dashed) contours of the objective ability distribution $N(\mathbf{0}, \mathbf{\Sigma}_r)$ (gray, mean marked by a circle) and of the estimated subjective prior $N(\boldsymbol{b}_r, e^{h_r}\mathbf{\Sigma}_r)$ (color, mean marked by a diamond), in the plane of the manual trades and managerial and professional clusters. Axes are log-wage ability units; all panels share the same axis limits. Each panel reports the type's population share π_r and the estimated variance scale e^{h_r}.

5.3 Model vs Data

While the estimation is still work in progress, the current parametrization qualitatively fits the targeted moments. In Figure 5(a) I report the model-implied employment shares across occupations and the switching hazard rate for the simulated model versus the data. I also show that the wage level and profile are matched — not perfectly, but in terms of gradient and increase over the lifecycle the model does a good job. The wage profile and wage growth over the lifecycle are not targeted moments. Importantly, the model can match these wage patterns only if it generates a "realistic" occupational pattern of which workers sort into which occupations.

Figure 5(b) shows the same moments for the different types. Occupational mobility, especially early in the career, is driven almost entirely by two types — the high-ability types, whose perceived uncertainty is still unresolved. These two types switch occupations at entry at rates of 0.47 and 0.53, while the lowest type never switches at entry and the second type barely does (0.03). However, the lowest type shows some occupational mobility after a few years in the

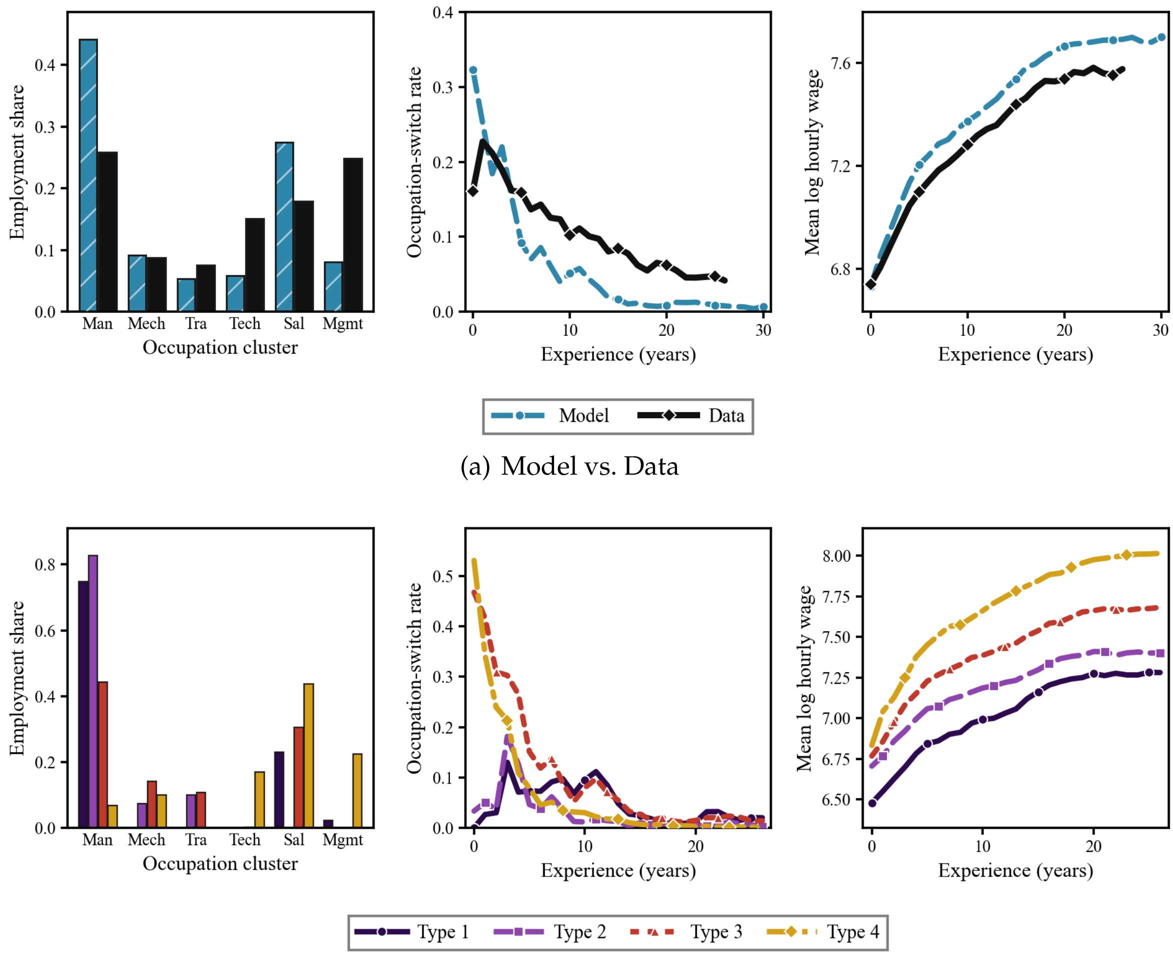

(a) Model vs. Data

(b) Model-simulated outcomes by worker type

Figure 5: Model fit and simulated moments by worker type. In each panel row: worker-year employment shares by occupation cluster (left), annual occupation-switching hazard by years of experience (middle), and mean log hourly wage by years of experience (right). Panel (a) compares model and data: solid bars and solid black lines are the data, hatched bars and dashed blue lines the model; data series are truncated at 26 years of experience, and wage moments are not targeted in the estimation. Panel (b) reports the model-simulated series disaggregated by worker type, ordered by ability (Type 1 lowest); series are drawn up to 26 years of experience for comparability with panel (a).

labor market. These moves are most likely driven by the accumulation of experience that is fully transferable across occupations: while accumulating experience, at some point another occupation may come to offer higher returns to that experience, and the worker moves there. The allocation across occupations follows the same logic: the managerial and professional cluster is populated almost exclusively by the top type (which places 23% of its workers there) and the three lower types are funneled into manual trades (employment shares of 0.44–0.83). The simulated wage profiles are also rank-ordered in type ability at entry, driven by the higher $\alpha_j + \gamma_{j,r}$. The highest-type workers have the steepest and largest wage increase, as they also sort into the occupations with the highest wage growth over the lifecycle, while Types 1 and 3 have a similar lifecycle profile. Although Type 3 workers move across occupations more than Type 1 workers, their moves reallocate them among the manual, sales and transport clusters, which offer relatively low returns to occupational tenure.

6 Experimentation as Compensating Differentials

Beyond wages and human capital accumulation, occupations differ in how much they let workers learn about their own abilities. In the model presented above, mean beliefs about unobserved abilities in a given occupation directly affect the worker's occupational choice, since workers compute the expected wage before sorting into an occupation. Because wages are log-linear, under log utility the expected flow utility is linear in mean beliefs. Given this functional form, belief dispersion does not directly affect the per-period returns to working in a given occupation. However, choosing an occupation affects mean beliefs about ability in that occupation, and potentially in other occupations as well, depending on the correlation structure and on belief dispersion in other occupations. Therefore, through the continuation value, beliefs enter the worker problem non-linearly, and this non-linearity makes belief dispersion and uncertainty matter. For a type-r worker in state $x = (\tilde{\boldsymbol{\xi}}, \boldsymbol{\tau})$, I define the experimentation value of occupation j as the part of its continuation value that is generated by the wage signal the occupation produces, that is, the change in the worker's continuation value from observing the current signal:

$$EV^{\text{info}}_{j,r}(x) = \beta\Big\{\mathbb{E}_s\big[\,V^r\big(\tilde{\boldsymbol{\xi}}^+_j(s),\, \boldsymbol{\tau}+\boldsymbol{e}_j\big)\big] - V^r\big(\tilde{\boldsymbol{\xi}},\, \boldsymbol{\tau}+\boldsymbol{e}_j\big)\Big\}, \tag{18}$$

where V^r is the value function of the dynamic program (10), $\tilde{\boldsymbol{\xi}}^+_j(s)$ is the belief mean after the belief update induced by the signal, and the expectation over the signal is taken under the worker's subjective belief distribution. Both terms keep the human-capital state $\boldsymbol{\tau}+\boldsymbol{e}_j$ fixed, so $EV^{\text{info}}_{j,r}$ isolates the value of the information the occupation delivers. Because the flow payoff is measured in log-wage units, $EV^{\text{info}}_{j,r}$ can be interpreted as an information compensating differential. In other words, it measures the log-wage amount a worker would be willing to give up in order to work in an occupation that provides valuable information about her ability and therefore the possibility of discovering that she is particularly productive in that occupation. An occupation whose experimentation value equals 0.05 log points is one the worker would enter at a five-percent wage discount relative to an occupation not providing any information, and the term enters the choice comparison in (11) additively, exactly like a wage premium. With Kalman updating, the signal moves the belief mean along a single known direction: writing $V = V^s_r(\boldsymbol{\tau})$ for the subjective posterior covariance (7), the updated mean is $\tilde{\boldsymbol{\xi}}^+_j(z) = \tilde{\boldsymbol{\xi}} + \boldsymbol{u}_j\, z$, with $z \sim \mathcal{N}(0,1)$ the standardized signal innovation and with the belief-updating vector

$$\boldsymbol{u}_j = \frac{V\boldsymbol{e}_j}{\sqrt{\boldsymbol{e}_j' V \boldsymbol{e}_j + \varsigma^2_{\varepsilon j}}}, \tag{19}$$

so that the expectation in (18) reduces to a one-dimensional Gaussian integral along this vector. The definition of the belief-updating vector clearly highlights what drives the value of experimentation in this model. The experimentation value is the option value of learning about ability in different occupations. Wage signals do not change beliefs on average — in fact, the expected value of the posterior is the prior itself — but they spread future mean beliefs around their current value. When the continuation value is convex in mean beliefs, this dispersion is valuable: good signals are informative about a productive match, while bad signals are less costly, as the worker is insured against bad draws — that is, she can switch occupation. The belief-updating vector in (19) determines how strongly occupation j moves beliefs, so more informative occupations generate a larger experimentation value.

Because abilities are correlated across occupations, one wage signal can move beliefs not only about the occupation where the signal is observed, but also about other occupations. This is reflected in the belief-updating vector $\boldsymbol{u}_j$, in which each component u_{jk} is the update of the mean belief in dimension k conditional on choosing occupation j. For example, u_{jj} denotes the revision of the mean belief about ability in occupation j induced by choosing occupation j. Therefore, the own-learning component of the experimentation value is defined exactly as in (18), except that the worker revises only the belief about the occupation she works in, while the spillover component is the remaining change in the continuation value, induced by the change in beliefs along the other coordinates:

$$EV_{j,r}^{\text{own}}(x) = \beta\Big\{\mathbb{E}_z\big[V^r\big(\tilde{\boldsymbol{\xi}} + u_{jj}\,\boldsymbol{e}_j\,z,\,\boldsymbol{\tau}+\boldsymbol{e}_j\big)\big] - V^r\big(\tilde{\boldsymbol{\xi}},\,\boldsymbol{\tau}+\boldsymbol{e}_j\big)\Big\}, \qquad EV_{j,r}^{\text{spill}} = EV_{j,r}^{\text{info}} - EV_{j,r}^{\text{own}}. \tag{20}$$

The own-learning component captures the part of the experimentation value coming from learning about ability in occupation j itself, while the spillover component captures the additional value coming from what the same signal reveals about abilities in other occupations. Both $EV_{j,r}^{\text{info}}$ and $EV_{j,r}^{\text{own}}$ are non-negative by the convexity argument in Appendix B.1, while the spillover component can be negative. Learning spillovers across occupations can be negative because learning about other occupations can reduce the incremental option value of learning about the current occupation. When occupations are substitutes in the continuation value, belief revisions that move several occupations together can compress differences in their relative attractiveness, making the signal less useful for future sorting. Importantly, total experimentation value remains non-negative; only the spillover component can be negative.

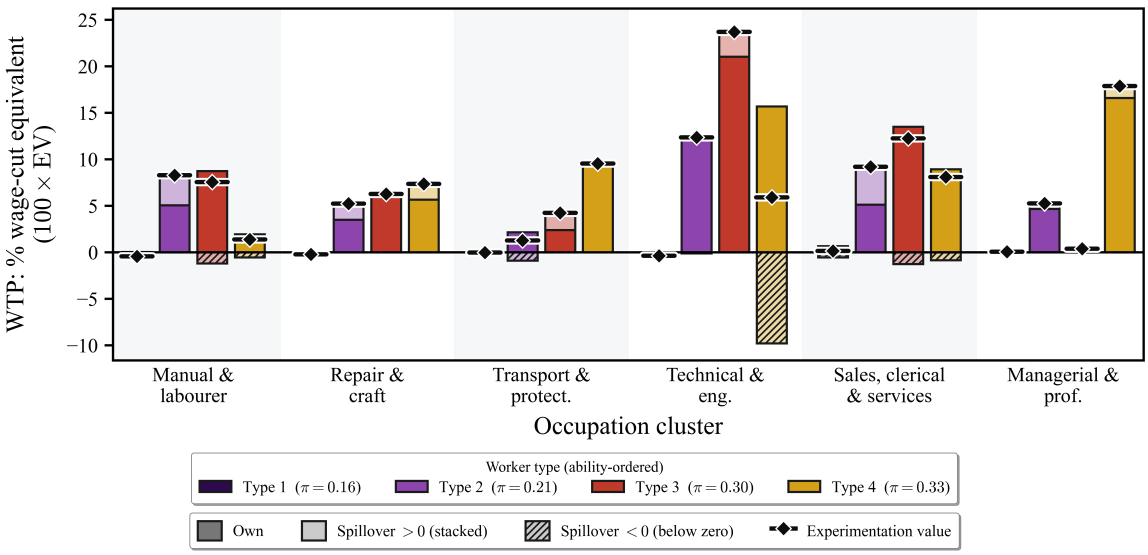

Figure 6: Own–spillover decomposition of the experimentation value at labor-market entry, by type and occupation cluster. For each occupation cluster and type, the bars show the own-learning and spillover components of the experimentation value in log points ($100 \times EV^{\text{info}}$), read as the percent wage cut the worker would accept in that occupation in exchange for its information; the diamond marks their sum, the total experimentation value. Types are ordered by ability (Type 1 lowest).

In Figure 6 I report the estimated experimentation value together with its two components, given the initial type-specific prior beliefs. Importantly, these components refer to the state before labor-market entry, before the worker chooses any occupation, and can be computed exactly for each type and occupation using all the estimated parameters and the approximated value function, without simulating the model.[5]

I find that the experimentation value across occupations and types is quite heterogeneous and depends strongly on the estimated type-specific prior beliefs. Type 1 workers, who have very low belief dispersion, extract practically no value from sorting into occupations in which they might have a good draw; that is, they experiment very little. In fact, their probability of switching occupation right after entering the labor market — when most of the learning about ability takes place — is almost zero. The other worker types, on the other hand, have heterogeneous returns to experimenting in different occupations. For most type–occupation pairs, the learning spillovers are positive or small in magnitude; the only sizable negative spillover is for Type 4 workers in the technical and engineering cluster.

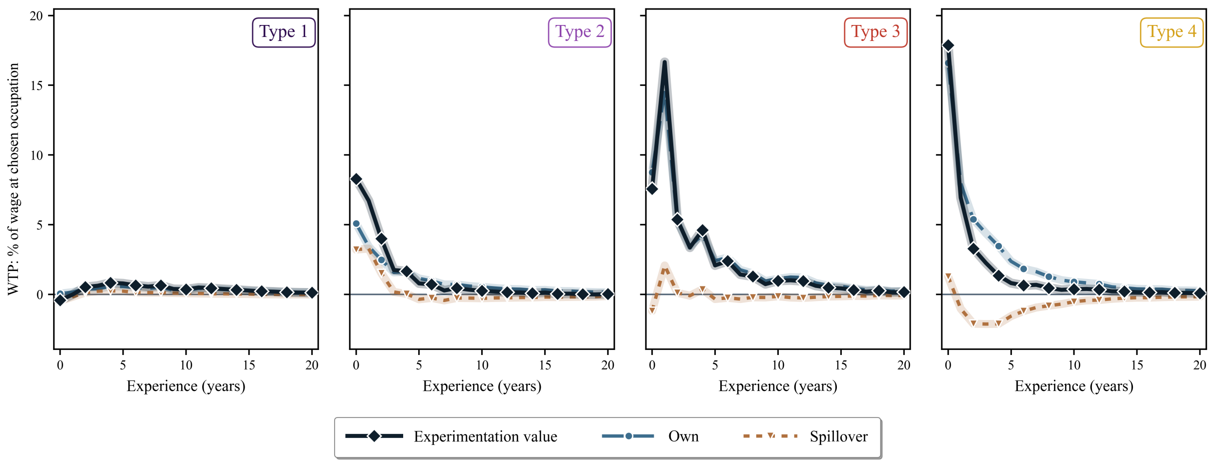

Figure 7: Experimentation value over the lifecycle as a willingness to pay, by type. One panel per type, ordered by ability (Type 1 lowest). The horizontal axis is years of experience; the vertical axis is the experimentation value at the worker's chosen occupation in log points ($100 \times EV^{\text{info}}$), read as the percent wage cut the worker would accept in that occupation in exchange for its information. Each panel plots the total experimentation value and its own-learning and spillover components; values are averages over simulated workers within each (type, experience) cell, evaluated at each worker's chosen occupation.

In Figure 7, I report the experimentation value of the simulated occupational choices over the lifecycle. The importance of the experimentation value declines over the lifecycle. Intuitively, belief updating reduces belief dispersion over time, and the experimentation motive fades as the probability of a large positive shock to mean beliefs goes down. It is almost always zero for

[5] At any state $x = (\tilde{\boldsymbol{\xi}}, \boldsymbol{\tau})$, I evaluate the belief-updating vector (19) at the estimated parameters and replace the value function with its sparse-grid interpolant $\widehat{V}^r$, so the expectation reduces to a univariate Gauss–Hermite sum,

$$\widehat{EV}_{j,r}(x) = \beta\Big[\pi^{-1/2}\sum_{g=1}^{G} w_g\,\widehat{V}^r\big(\tilde{\boldsymbol{\xi}} + \sqrt{2}\,\zeta_g\,\hat{\boldsymbol{u}}_j,\ \boldsymbol{\tau} + \boldsymbol{e}_j\big) - \widehat{V}^r\big(\tilde{\boldsymbol{\xi}},\ \boldsymbol{\tau} + \boldsymbol{e}_j\big)\Big],$$

with (ζ_g, w_g) the Gauss–Hermite nodes and weights; the own component replaces $\hat{\boldsymbol{u}}_j$ with $\hat{u}_{jj}\boldsymbol{e}_j$. The lifecycle profiles below report $\widehat{EV}_{d_{it},r}(x_{it})$ at the occupation d_{it} chosen by worker i at experience t, averaged over the simulated workers' careers.

Type 1 workers, and increasing in the types' average ability: higher types extract more value from experimenting. For Type 4 workers, the spillover effects are actually negative and partly offset the effect of learning about ability in the chosen occupation.

7 Conclusion

In this paper, I show that in a dynamic occupational choice model where workers sort based on their expected income an experimentation motive arises naturally through the continuation value. Early in the career, when uncertainty about abilities is higher, workers value positive ability realizations more than negative ones. I then estimate the model on NLSY79 panel data to quantify this experimentation motive, allowing for unobserved heterogeneity that is identified through ability measurements that are independent of labor-market outcomes. This approach follows Arcidiacono et al. (2025), but importantly I use observed occupational choice data to identify the subjective beliefs of each type, therefore relaxing the rational-expectations assumption often imposed in previous work in the literature.

I show that experimenting acts as a compensating differential, and the willingness to pay for trying out occupations in which ability is more uncertain is decreasing over the lifecycle and heterogeneous across types. High-ability workers are those who have a higher willingness to pay for uncertainty resolution and are those that experiment more over the lifecycle.

References

Elizabeth S. Allman, Catherine Matias, and John A. Rhodes. Identifiability of parameters in latent structure models with many observed variables. *The Annals of Statistics*, 37(6A):3099–3132, 2009.

Kate Antonovics and Limor Golan. Experimentation and job choice. *Journal of Labor Economics*, 30(2):333–366, 2012.

Peter Arcidiacono, Esteban Aucejo, Arnaud Maurel, and Tyler Ransom. College attrition and the dynamics of information revelation. *Journal of Political Economy*, 133(1):53–110, 2025.

Johannes Brumm and Simon Scheidegger. Using adaptive sparse grids to solve high-dimensional dynamic models. *Econometrica*, 85(5):1575–1612, 2017.

Jackson Bunting, Paul Diegert, and Arnaud Maurel. Heterogeneity, uncertainty and learning: Semiparametric identification and estimation. Technical report, National Bureau of Economic Research, 2024.

Piero De Dominicis. Multidimensional skills, learning and occupational sorting. *Learning and Occupational Sorting (April 07, 2025)*, 2025.

Fatih Guvenen, Burhan Kuruscu, Satoshi Tanaka, and David Wiczer. Multidimensional skill mismatch. *American Economic Journal: Macroeconomics*, 12(1):210–44, 2020.

Boyan Jovanovic. Job matching and the theory of turnover. *Journal of Political Economy*, 87(5): 972–990, 1979.

Robert A Miller. Job matching and occupational choice. *Journal of Political economy*, 92(6):1086–1120, 1984.

Theodore Papageorgiou. Learning your comparative advantages. *Review of Economic Studies*, 81(3):1263–1295, 2014.

Elena Pastorino. Careers in firms: The role of learning about ability and human capital acquisition. *Journal of Political Economy*, 132(6):1994–2073, 2024.

A Appendix: Data

A.1 Sample composition and summary statistics

Table 2 reports summary statistics for the estimation sample. The first panel reports average labor-market characteristics at the worker-year level; the second shows lifecycle characteristics averaged across workers; the third reports raw number-correct scores on the seven ASVAB subtests entering the measurement system (in estimation these are residualized by age and standardized); the fourth reports the demographics of the sample in terms of race and years of schooling.

A.2 Occupation clustering

Clusters are obtained by using a k-means algorithm with $J = 6$ centroids on the occupation-level means of all 120 standardized O*NET descriptors. Each occupation is weighted by its worker-year employment mass in the panel.

The partition explains 47.5% of the employment-weighted total task-content variance across occupations (between-cluster share of total sum of squares in the standardized descriptor space).

Occupation-cluster tenure is computed from the raw occupation histories under the partition: prior observed worker-years in the cluster, starting at zero, with no reset on exit and re-entry.

B Appendix: Model

B.1 Convexity in Beliefs and Non-negativity of EV^{info}

Fix a worker type r and, when unambiguous, suppress the type superscript to simplify notation. Recall the state $x = (\tilde{\boldsymbol{\xi}}, \boldsymbol{\tau})$. The wage signal $s = \boldsymbol{e}_j'\boldsymbol{\xi} + \varepsilon_j$ moves the worker's beliefs by the Kalman update in (10), and it does two things: it moves her best guess about her abilities, and it lowers her perceived uncertainty about them. The second effect is carried by the tenure

	Mean	SD	N
Wages, age and tenure			
Log hourly wage	7.284	0.687	43,342
Age	33.861	8.472	43,342
Occupational tenure	6.585	6.166	43,342
Career history (worker-level)			
Distinct three-digit occupations held	4.050	2.327	1,871
Person-years per worker	23.17	7.21	1,871
ASVAB subtests (raw number-correct)			
Arithmetic Reasoning (30 items)	18.694	7.584	1,871
Word Knowledge (35 items)	25.733	7.682	1,871
Paragraph Comprehension (15 items)	10.603	3.482	1,871
Mathematics Knowledge (25 items)	14.252	6.722	1,871
Mechanical Comprehension (25 items)	15.759	5.342	1,871
Auto & Shop Information (25 items)	16.304	5.269	1,871
Electronics Information (20 items)	12.561	4.217	1,871
Demographics (worker-level)			
Black	0.113	0.316	1,871
Hispanic	0.067	0.251	1,871
Years of schooling	13.561	2.548	1,871

Table 2: Summary statistics for the estimation sample. Source: NLSY79 + O*NET; sample construction follows criteria similar to Guvenen et al. (2020).

Cluster	Top occupations by worker-years	n_{occ}	Share
0	Machine operators n.e.c., carpenters, construction laborers, janitors	89	0.259
1	Automobile mechanics, mechanics and repairers n.e.c., electricians, machinists	23	0.088
2	Truck and delivery drivers, guards and watchmen, airplane pilots	11	0.075
3	Production supervisors, software developers, systems analysts, engineers	47	0.151
4	Salespersons n.e.c., cooks, police and detectives, shipping clerks	79	0.179
5	Managers and administrators n.e.c., accountants, financial managers, lawyers	40	0.249

Table 3: Occupation clusters on the full O*NET task space: representative occupations (by worker-year employment), number of three-digit occupation codes, and worker-year employment shares.

increment — by the Sherman–Morrison identity, $V_r^s(\boldsymbol{\tau} + \boldsymbol{e}_j)$ is exactly the post-signal covariance — and is common to both terms of (18), which therefore prices the movement of the belief mean alone. This appendix shows that this price is non-negative.

Using the standardized signal innovation $z \sim \mathcal{N}(0, 1)$, the posterior mean after choosing occupation j can be written as

$$\tilde{\boldsymbol{\xi}}_j^+ = \tilde{\boldsymbol{\xi}} + \boldsymbol{u}_j(\boldsymbol{\tau})\, z,$$

where $\boldsymbol{u}_j(\boldsymbol{\tau})$ is the belief-updating vector in (19). Crucially, the updating vector depends on the state only through the tenure vector: beliefs move along a direction that is known before the signal is realized and does not depend on the current mean belief $\tilde{\boldsymbol{\xi}}$. Collecting in $a_j(\boldsymbol{\tau})$ all the flow-payoff terms that do not depend on beliefs, the stationary Bellman equation (10) can be written as

$$V(\tilde{\boldsymbol{\xi}}, \boldsymbol{\tau}) = \max_{j \in \mathcal{O}} \Big\{ a_j(\boldsymbol{\tau}) + \tilde{\xi}_j + \beta \, \mathbb{E}_z \big[V\big(\tilde{\boldsymbol{\xi}} + \boldsymbol{u}_j(\boldsymbol{\tau})\, z,\, \boldsymbol{\tau} + \boldsymbol{e}_j\big) \big] \Big\}. \tag{21}$$

Proposition 1 (Convexity in mean beliefs) *For every type r and every tenure state $\boldsymbol{\tau}$, the value function*

$$\tilde{\boldsymbol{\xi}} \longmapsto V^r(\tilde{\boldsymbol{\xi}}, \boldsymbol{\tau})$$

is convex.

Proof *Let T denote the Bellman operator defined by the right-hand side of* (21)*, and suppose that the candidate continuation function $f(\cdot, \boldsymbol{\tau}')$ is convex in mean beliefs for every tenure state $\boldsymbol{\tau}'$. For every occupation j and every realization of z, the mapping*

$$\tilde{\boldsymbol{\xi}} \longmapsto \tilde{\boldsymbol{\xi}} + \boldsymbol{u}_j(\boldsymbol{\tau})\, z$$

is affine, and composing a convex function with an affine map preserves convexity, so $f\big(\tilde{\boldsymbol{\xi}} + \boldsymbol{u}_j(\boldsymbol{\tau})\, z,\, \boldsymbol{\tau} + \boldsymbol{e}_j\big)$ is convex in $\tilde{\boldsymbol{\xi}}$ for every z. The expectation operator over z preserves convexity, and so does adding the current payoff $a_j(\boldsymbol{\tau}) + \tilde{\xi}_j$, which is affine in $\tilde{\boldsymbol{\xi}}$. Hence every choice-specific value is convex in mean beliefs, and since the pointwise maximum of convex functions is convex, T maps functions convex in mean beliefs into functions convex in mean beliefs.

Now initialize value iteration at $V_0 \equiv 0$, which is convex, and let $V_{n+1} = TV_n$. Because T preserves convexity, $V_n(\cdot, \boldsymbol{\tau})$ is convex for every n, and because $\beta \in (0, 1)$, value iteration converges pointwise to the Bellman fixed point V. A pointwise finite limit of convex functions is convex, hence $V(\cdot, \boldsymbol{\tau})$ is convex.

Corollary 1 (Non-negativity of the experimentation value) *For every state x, every occupation j, and every type r, $EV^{info}_{j,r}(x) \geq 0$ and $EV^{own}_{j,r}(x) \geq 0$.*

Proof *The signal innovation has mean zero, so $\mathbb{E}_z\big[\tilde{\boldsymbol{\xi}} + \boldsymbol{u}_j(\boldsymbol{\tau})\, z\big] = \tilde{\boldsymbol{\xi}}$: the expected posterior mean is the prior mean. Jensen's inequality applied to the convex function $V^r(\cdot, \boldsymbol{\tau} + \boldsymbol{e}_j)$ then gives*

$$\mathbb{E}_z \big[V^r\big(\tilde{\boldsymbol{\xi}} + \boldsymbol{u}_j(\boldsymbol{\tau})\, z,\, \boldsymbol{\tau} + \boldsymbol{e}_j\big) \big] \geq V^r\big(\tilde{\boldsymbol{\xi}},\, \boldsymbol{\tau} + \boldsymbol{e}_j\big),$$

so $EV^{info}_{j,r}(x) \geq 0$ in (18)*. The same argument applies verbatim with $\boldsymbol{u}_j$ replaced by $u_{jj}\, \boldsymbol{e}_j$ — the innovation remains mean zero along the shorter ray — so $EV^{own}_{j,r}(x) \geq 0$ in* (20)*.*

No such bound holds for the residual $EV^{\text{spill}}_{j,r} = EV^{\text{info}}_{j,r} - EV^{\text{own}}_{j,r}$: convexity does not order the full-ray gain against the own-coordinate gain, and cross-coordinate complementarities or substitutabilities in the continuation value can make the spillover term negative.

C Appendix: Estimation Results

C.1 Stage 1 Estimation Results: EM Parameter Estimates

Table 4 reports the full Stage-1 point estimates: the per-cluster wage-equation parameters (Panel A), the saturated type premia (Panel B), selected step coefficients of the seniority functions F_j and H_j (Panels C and D), and the type shares and measurement-block parameters (Panel E); the type-specific ability covariance matrices are reported in Table 6 and the implied correlations in Figure 8 below.

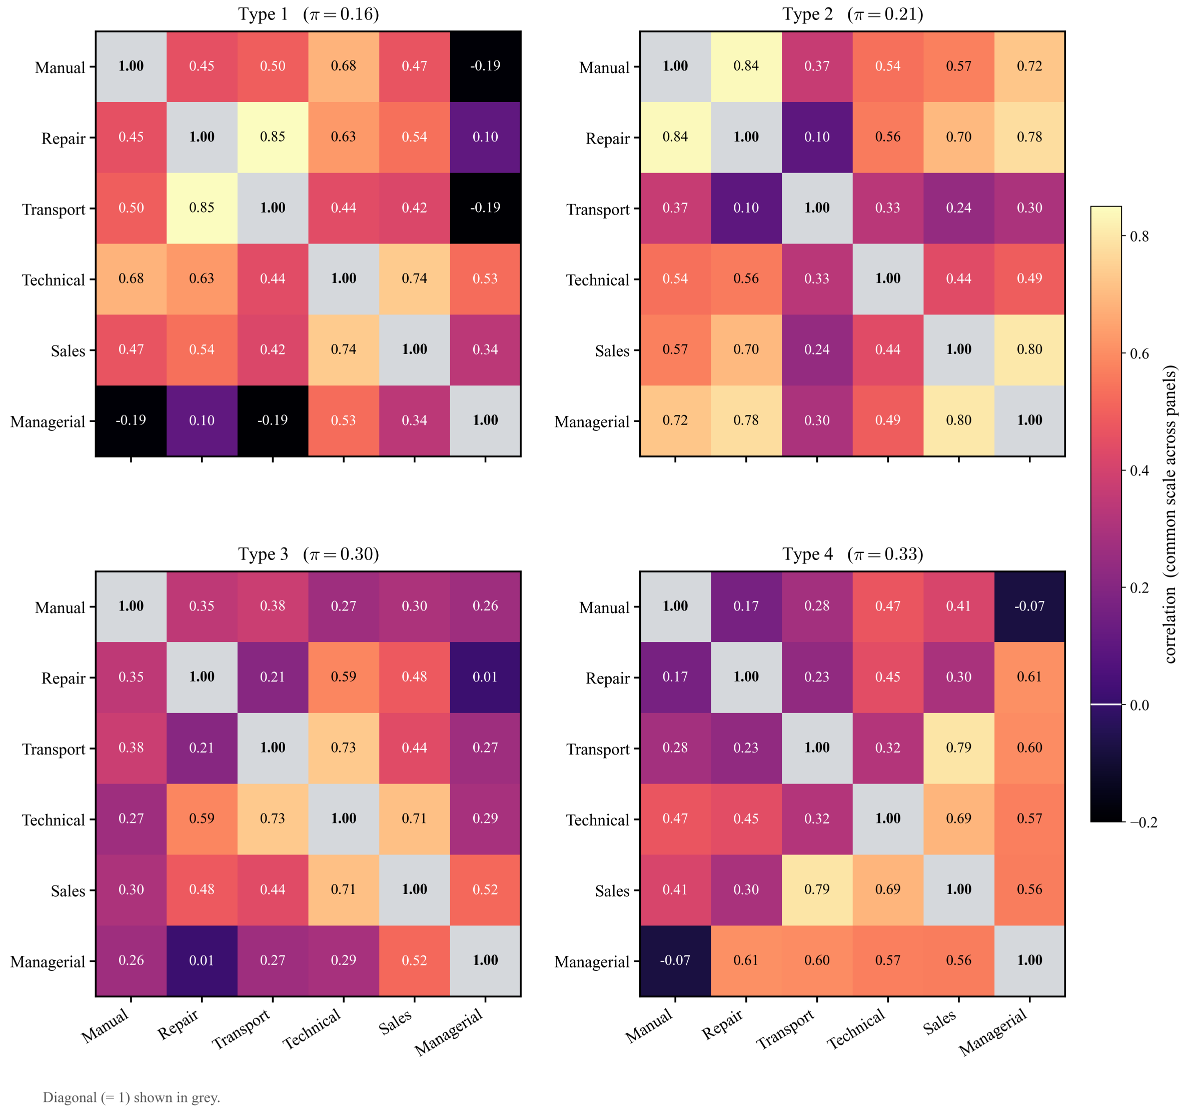

Figure 8: *Type-specific correlation matrices of occupation-specific abilities*, implied by the estimated covariances. One panel per type, ordered by ability (Type 1, lowest, top left, to Type 4, highest, bottom right); each panel label reports the type share π_r. Color is a sequential map over the observed range of correlations on a scale shared by the four panels; the unit diagonal is masked in gray; occupation clusters are labeled on the axes. The underlying covariance matrices are reported in Appendix Table 6.

Table 4: Stage 1: Expectation–Maximization Estimates

Panel A: Wage Equation by Occupation Cluster

Cluster	α_j	$\varsigma^2_{\varepsilon j}$
0 — manual, laborer	6.695	0.111
1 — skilled repair, craft	6.670	0.112
2 — transport, protective	6.646	0.125
3 — technical, engineering	6.612	0.099
4 — sales, clerical, services	6.502	0.142
5 — managerial, professional	6.609	0.320

Panel B: Type Premia $\gamma_{j,r}$ (relative to Type 1)

Cluster	Type 1	Type 2	Type 3	Type 4
0 — manual, laborer	0	0.223	0.311	0.336
1 — skilled repair, craft	0	0.137	0.260	0.341
2 — transport, protective	0	0.128	0.269	0.314
3 — technical, engineering	0	0.133	0.315	0.578
4 — sales, clerical, services	0	0.178	0.379	0.511
5 — managerial, professional	0	0.194	0.404	0.642

Panel C: Tenure Steps $F_j(\tau)$, Selected Bins

Cluster	$\tau = 1$	2	3	5	7	10–11	14–15	18+
0 — manual, laborer	0.058	0.079	0.100	0.167	0.168	0.217	0.235	0.242
1 — skilled repair, craft	0.025	0.075	0.116	0.103	0.089	0.157	0.195	0.133
2 — transport, protective	0.052	0.073	0.089	0.087	0.173	0.206	0.198	0.155
3 — technical, engineering	0.044	0.112	0.150	0.190	0.223	0.258	0.315	0.342
4 — sales, clerical, services	0.017	0.027	0.052	0.100	0.102	0.163	0.188	0.237
5 — managerial, professional	0.022	0.068	0.110	0.140	0.158	0.204	0.354	0.422

Panel D: General-Experience Steps $H_j(E)$, Selected Bins

Cluster	$E = 1$	2	3	5	10–11	14–15	20–21	24+
0 — manual, laborer	0.009	0.064	0.117	0.158	0.221	0.310	0.406	0.404
1 — skilled repair, craft	0.049	0.078	0.157	0.254	0.308	0.362	0.518	0.534
2 — transport, protective	−0.003	0.049	0.079	0.132	0.254	0.268	0.386	0.380
3 — technical, engineering	0.017	0.070	0.116	0.246	0.417	0.530	0.621	0.639
4 — sales, clerical, services	0.029	0.105	0.151	0.293	0.406	0.533	0.619	0.647
5 — managerial, professional	0.079	0.101	0.162	0.268	0.449	0.458	0.472	0.542

Notes: Types are ordered by ability (Type 1 lowest, Type 4 highest). Panel A reports the occupation fixed effect and the variances of the idiosyncratic shocks. Panel B reports the occupation-by-type premia normalized in each cluster with respect to the lowest-ability type. Panels C and D report the estimated stepwise return functions for occupational tenure and general experience; bins are defined yearly up to nine years, afterwards each bin groups two years of tenure and experience until 18 years of occupational tenure and 24 years of general experience. Beyond these values, returns do not change.

Table 5: Stage 1: Expectation–Maximization Estimates, continued

Panel E: Type Shares and Measurement Block

	Type 1	Type 2	Type 3	Type 4	
Type share π_r	0.158	0.212	0.295	0.335	
ASVAB subtest	μ_{1p}	μ_{2p}	μ_{3p}	μ_{4p}	$\sigma^2_{\nu p}$
Word Knowledge	−1.66	−0.44	0.31	0.83	0.248
Paragraph Comprehension	−1.52	−0.50	0.31	0.81	0.327
Arithmetic Reasoning	−1.26	−0.83	0.09	1.06	0.220
Mathematics Knowledge	−1.14	−0.83	−0.06	1.12	0.214
Mechanical Comprehension	−1.28	−0.55	0.17	0.84	0.413
Auto & Shop Information	−1.28	−0.27	0.26	0.61	0.560
Electronics Information	−1.38	−0.51	0.19	0.84	0.379

Notes: Panel E reports the mixture probabilities and the measurement block: the unrestricted type means on the seven age-residualized, standardized ASVAB subtests and the diagonal measurement noise variances. Types are ordered by ability (Type 1 lowest, Type 4 highest). The type-specific ability covariance matrices are in Table 6; the implied correlations are displayed in Figure 8. Point estimates; bootstrap standard errors are pending.

$$\mathbf{\Sigma}_1 = \begin{pmatrix} 0.16 & 0.07 & 0.06 & 0.08 & 0.08 & -0.03 \\ 0.07 & 0.16 & 0.11 & 0.08 & 0.09 & 0.02 \\ 0.06 & 0.11 & 0.10 & 0.04 & 0.05 & -0.03 \\ 0.08 & 0.08 & 0.04 & 0.09 & 0.09 & 0.07 \\ 0.08 & 0.09 & 0.05 & 0.09 & 0.17 & 0.06 \\ -0.03 & 0.02 & -0.03 & 0.07 & 0.06 & 0.18 \end{pmatrix} \quad \mathbf{\Sigma}_2 = \begin{pmatrix} 0.09 & 0.06 & 0.03 & 0.05 & 0.06 & 0.08 \\ 0.06 & 0.05 & 0.01 & 0.04 & 0.05 & 0.07 \\ 0.03 & 0.01 & 0.09 & 0.03 & 0.02 & 0.03 \\ 0.05 & 0.04 & 0.03 & 0.11 & 0.05 & 0.06 \\ 0.06 & 0.05 & 0.02 & 0.05 & 0.10 & 0.10 \\ 0.08 & 0.07 & 0.03 & 0.06 & 0.10 & 0.14 \end{pmatrix}$$

$$\mathbf{\Sigma}_3 = \begin{pmatrix} 0.12 & 0.04 & 0.04 & 0.04 & 0.04 & 0.04 \\ 0.04 & 0.09 & 0.02 & 0.08 & 0.06 & 0.00 \\ 0.04 & 0.02 & 0.11 & 0.11 & 0.06 & 0.04 \\ 0.04 & 0.08 & 0.11 & 0.20 & 0.13 & 0.06 \\ 0.04 & 0.06 & 0.06 & 0.13 & 0.16 & 0.09 \\ 0.04 & 0.00 & 0.04 & 0.06 & 0.09 & 0.19 \end{pmatrix} \quad \mathbf{\Sigma}_4 = \begin{pmatrix} 0.15 & 0.03 & 0.05 & 0.08 & 0.07 & -0.02 \\ 0.03 & 0.17 & 0.04 & 0.08 & 0.05 & 0.15 \\ 0.05 & 0.04 & 0.21 & 0.06 & 0.16 & 0.16 \\ 0.08 & 0.08 & 0.06 & 0.19 & 0.13 & 0.14 \\ 0.07 & 0.05 & 0.16 & 0.13 & 0.19 & 0.14 \\ -0.02 & 0.15 & 0.16 & 0.14 & 0.14 & 0.34 \end{pmatrix}$$

Table 6: Type-specific ability covariance matrices $\mathbf{\Sigma}_r$, types ordered by ability (Type 1 lowest, Type 4 highest). Rows and columns index the occupation clusters $j = 0, \ldots, 5$ of Table 1: 0 manual/laborer, 1 skilled repair/craft, 2 transport/protective, 3 technical/engineering, 4 sales/clerical/services, 5 managerial/professional. The implied correlation matrices are displayed in Figure 8. Point estimates; bootstrap standard errors are pending.

	Manual	Mechanics	Transport	Technical	Sales	Management
Type 1	+0.162	−0.046	−0.048	−0.138	+0.193	−0.123
Type 2	+0.096	+0.043	+0.102	−0.153	−0.080	−0.008
Type 3	+0.175	+0.106	+0.162	−0.098	+0.149	−0.494
Type 4	+0.158	+0.073	+0.053	−0.223	+0.107	−0.169

Table 7: Estimated type-specific prior means $\boldsymbol{b}_r$ (perceived comparative advantage at labor-market entry), in log-wage ability units. Types are ordered by ability (Type 1 lowest); columns are the occupation clusters of Table 1. Each row is normalized to sum to zero. Preliminary point estimates; bootstrap standard errors are pending.

C.2 Functional Form of Occupational Tenure and General Experience

Figure 9 compares the estimated return functions with a smooth quadratic specification, $F_j(\tau) = \phi_{1j}\tau + \phi_{2j}\tau^2$ and $H_j(E) = \chi_{1j}E + \chi_{2j}E^2$, estimated on the same sample. In the main specification of the model, I use stepwise functions to avoid extreme curvature that would be implied by quadratic functions for occupations in which there are not many observations in the right tail of the tenure/experience distributions. For example, the return to occupational tenure turns negative too quickly for the cluster “Skilled repair and craft”.

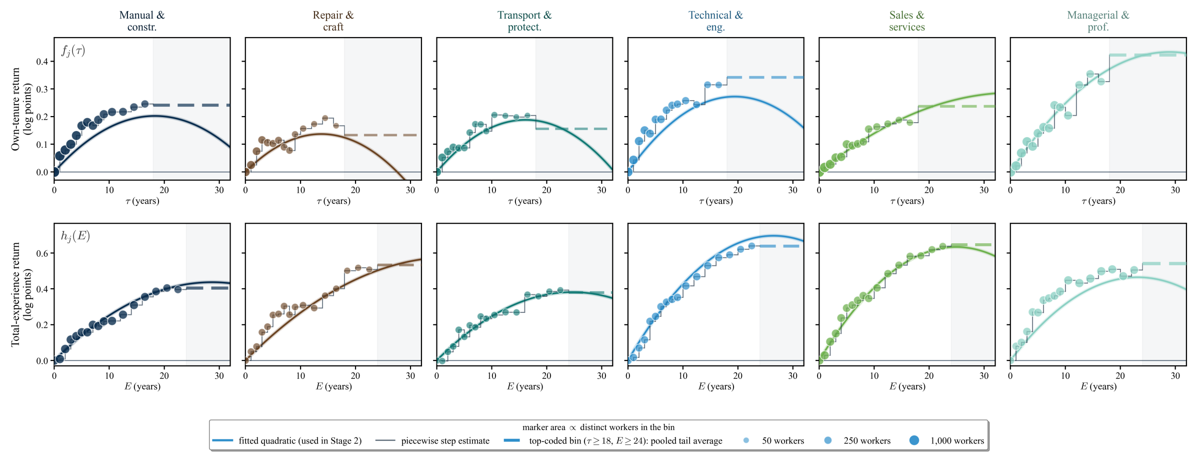

Figure 9: Stepwise versus quadratic returns to occupational tenure and general experience for each cluster. The quadratic specification attains a log-likelihood of −38,167.82 against −38,045.19 for the stepwise specification, but the stepwise specification has 162 additional parameters.

D Appendix: EM Algorithm — Computational Details

This appendix states the explicit E-step and M-step that are iterated in Stage 1 to estimate the wage equation parameters. In what follows, n indexes worker-year observations, $i(n)$ the corresponding worker, and d_n the occupation cluster of observation n; the deterministic wage mean at type r is

$$\mu_n(r) = \alpha_{d_n} + \gamma_{d_n,r} + F_{d_n}(\tau_n) + H_{d_n}(E_n),$$

with $E_n = \mathbf{1}'\boldsymbol{\tau}_n$ the worker's general experience at observation n, and F_j, H_j evaluated at the estimated step coefficients for occupational tenure and general experience returns.

D.1 E-step

Stack worker i's wages into $\boldsymbol{w}_i \in \mathbb{R}^{T_i}$, the unit-loading design into $\Lambda_i \in \mathbb{R}^{T_i \times J}$ with rows $\boldsymbol{e}'_{d_{it}}$, and $\Omega_i = \mathrm{diag}(\varsigma^2_{\varepsilon,d_{it}})$. Gaussian conjugacy delivers the type-conditional posterior of $\boldsymbol{\xi}_i$,

$$V_{ir} = \left(\boldsymbol{\Sigma}_r^{-1} + \Lambda_i'\Omega_i^{-1}\Lambda_i\right)^{-1} = \left(\boldsymbol{\Sigma}_r^{-1} + \mathrm{diag}(\boldsymbol{n}_i \oslash \varsigma^2_\varepsilon)\right)^{-1}, \tag{22}$$

$$\boldsymbol{m}_{ir} = V_{ir}\,\Lambda_i'\Omega_i^{-1}\left(\boldsymbol{w}_i - \boldsymbol{\mu}_i(r)\right), \tag{23}$$

where $\boldsymbol{n}_i$ counts worker i's observations per cluster and $\oslash$ is elementwise division — the wage information matrix is diagonal because of the unit loadings, so the posterior precision is the type-specific prior precision plus a tenure-driven diagonal, mirroring (7). Unlike the common-covariance case, V_{ir} depends on r through $\boldsymbol{\Sigma}_r^{-1}$ and is computed and stored per (i, r) pair. The wage marginal likelihood $\boldsymbol{w}_i \mid r \sim \mathcal{N}(\boldsymbol{\mu}_i(r), \Lambda_i \boldsymbol{\Sigma}_r \Lambda_i' + \Omega_i)$ is evaluated as

$$\log L_i^{w,\mathrm{marg}}(r) = -\tfrac{1}{2}\Big[T_i \log(2\pi) + \textstyle\sum_t \log \varsigma^2_{\varepsilon, d_{it}} + \log\det \boldsymbol{\Sigma}_r - \log\det V_{ir} + \mathcal{Q}_i(r)\Big], \tag{24}$$

with $\mathcal{Q}_i(r) = \sum_t (w_{it} - \mu_{it}(r))^2 / \varsigma^2_{\varepsilon, d_{it}} - \boldsymbol{c}_{ir}' \boldsymbol{m}_{ir}$ and $\boldsymbol{c}_{ir} \equiv \Lambda_i' \Omega_i^{-1}(\boldsymbol{w}_i - \boldsymbol{\mu}_i(r))$. The measurement marginal is the product of seven univariate normal densities at $(\boldsymbol{\mu}_r, \boldsymbol{\Sigma}_\nu)$, and the type posterior ω_{ir} follows from (15) by log-sum-exp.

D.2 M-step

(M1) Mixture weights and ability covariances. The mixture weights for different types and the ability variance-covariance matrices are computed as follows:

$$\pi_r = \frac{1}{N}\sum_{i=1}^{N} \omega_{ir}, \qquad \boldsymbol{\Sigma}_r = \frac{\sum_{i=1}^{N} \omega_{ir}\left(V_{ir} + \boldsymbol{m}_{ir}\boldsymbol{m}_{ir}'\right)}{\sum_{i=1}^{N} \omega_{ir}}. \tag{25}$$

Importantly, unlike in De Dominicis (2025), $\boldsymbol{\Sigma}_r$ represents the posterior second moment within its own type.

(M2) Occupation-by-occupation wage block. For each cluster j, the coefficient vector (which includes the occupation fixed effect, the type premia, and the occupation-specific returns to occupational tenure and general experience)

$$\boldsymbol{\theta}_j \equiv \left(\alpha_j, \gamma_{j,2}, \ldots, \gamma_{j,R}, \{F_j(k)\}_k, \{H_j(k')\}_{k'}\right)$$

is estimated by type-weighted least squares, with the dependent variable equal to the observed wage net of the worker's posterior ability mean,

$$y_{nr} = w_n - [\boldsymbol{m}_{i(n),r}]_j, \qquad \boldsymbol{x}_{nr} = \left(1, \mathbf{1}\{r=2\}, \ldots, \mathbf{1}\{r=R\}, \boldsymbol{\iota}^F(\tau_n), \boldsymbol{\iota}^H(E_n)\right),$$

where $\boldsymbol{\iota}^F(\tau_n)$ and $\boldsymbol{\iota}^H(E_n)$ stack the tenure- and experience-bin indicators, with weights $\omega_{i(n),r}$ over all pairs (n, r) with $d_n = j$. Because the ability enters additively with unit loading, the posterior-variance term $[V_{ir}]_{jj}$ is additively separable from the coefficients in the Q-function and the updating equations do not require a posterior-variance correction; this correction term appears only in the residual variance,

$$\varsigma^2_{\varepsilon j} = \frac{\sum_{(n,r):\, d_n = j} \omega_{i(n),r}\left[\left(y_{nr} - \boldsymbol{\theta}_j' \boldsymbol{x}_{nr}\right)^2 + [V_{i(n),r}]_{jj}\right]}{\sum_{(n,r):\, d_n = j} \omega_{i(n),r}}, \tag{26}$$

where $[V_{ir}]_{jj}$ must be carried per type — the single place in the M-step where type-specificity of the posterior covariance changes the algebra relative to the common-covariance model.

(M3) Measurement block. The updates for the measurement system parameters are simply weighted means and residual moments,

$$\boldsymbol{\mu}_r = \frac{\sum_i \omega_{ir}\, \boldsymbol{y}_i}{\sum_i \omega_{ir}}, \qquad \sigma^2_{vp} = \frac{1}{N}\sum_{i=1}^{N}\sum_{r=1}^{R} \omega_{ir}\big(y_{ip} - \mu_{rp}\big)^2. \tag{27}$$

D.3 Convergence

The EM algorithm stops iterating the two steps upon convergence of the observed data log-likelihood, which occurs when $0 \le \ell^{(q)} - \ell^{(q-1)} < 10^{-7}$. Because of the presence of local optima, I initialize the EM algorithm from different initial guesses; the reported estimates attain the highest converged likelihood across starts.

E Appendix: Stage-2 Algorithm — Computational Details

E.1 State space and value-function approximation

The per-type continuous state is $(\tilde{\boldsymbol{\xi}}, \boldsymbol{\tau}) \in \mathbb{R}^{2J} = \mathbb{R}^{12}$; the subjective posterior covariance is not a state, being the deterministic function (7) of tenure, recomputed on the fly. Because the type is known to the worker, the R Bellman systems decouple and are solved in parallel. The fixed point of the max Bellman operator (17) is kinked at belief indifference points, which dictates the choice of approximation method.

E.2 Continuation-value expectation: one-dimensional quadrature

Conditional on choosing a given occupation j, the covariance and tenure transitions are deterministic and mean beliefs move along the single direction determined by $\boldsymbol{u}_j$ of (19). For any function G,

$$\mathbb{E}_s\big[G\big(\tilde{\boldsymbol{\xi}}_j^{+}(s)\big)\big] = \int_{-\infty}^{\infty} G\big(\tilde{\boldsymbol{\xi}} + \boldsymbol{u}_j z\big)\, \varphi(z)\, dz \approx \pi^{-1/2}\sum_{g=1}^{N_q} w_g\, G\big(\tilde{\boldsymbol{\xi}} + \sqrt{2}\,\boldsymbol{u}_j \zeta_g\big), \tag{28}$$

with nodes and weights $\{(\zeta_g, w_g)\}$ from the univariate Gauss–Hermite rule: the integration is one-dimensional regardless of J, and exact for polynomials in z up to degree $2N_q - 1$. The experimentation-value evaluations of Section 6 use this same primitive, so the reported EV^{info} is computed with the estimator's own integrand.

E.3 Inner loop, simulation, and outer loop

Given a candidate $\boldsymbol{\Theta}_3 = (\boldsymbol{b}_r, h_r)_{r=1}^{R}$, the inner loop solves the max Bellman fixed point (17) for each type (contraction modulus $\beta = 0.9$; warm-started across outer evaluations). The simulation step then generates career paths and computes the simulated moments entering the criterion (16). The outer loop updates $\mathbf{\Theta}_3$ to decrease the criterion.

E.4 Two-stage algorithm summary

Algorithm: Two-Stage Estimation (EM + SMM)

Inputs. Panel $\{(\boldsymbol{w}_i, \boldsymbol{y}_i, \boldsymbol{d}_i, \boldsymbol{\tau}_i)\}_{i=1}^N$; occupation-cluster partition from the full O*NET task space.
Outputs. Estimates $(\widehat{\mathbf{\Theta}}_1, \widehat{\mathbf{\Theta}}_2, \widehat{\mathbf{\Theta}}_3)$, type posteriors $\widehat{\omega}_{ir}$, belief moments $(\widehat{\boldsymbol{m}}_{ir}, \widehat{V}_{ir})$.

Stage 1 — EM on $(\mathbf{\Theta}_1, \mathbf{\Theta}_2)$ (Section 4.4.1, Appendix D).

1. **Initialize**: per-cluster OLS for $(\alpha_j, F_j, H_j, \varsigma^2_{\varepsilon j})$; $\mathbf{\Sigma}_r^{(0)}$ at the common-covariance optimum; uniform $\pi_r^{(0)} = 1/R$; k-means starts for $\boldsymbol{\mu}_r^{(0)}$.
2. **Repeat** until $0 \le \ell^{(q)} - \ell^{(q-1)} < 10^{-7}$:
   E-step: $(V_{ir}, \boldsymbol{m}_{ir})$ via (22)–(23); wage marginal likelihood (24);
   measurement marginal; type posteriors ω_{ir} by log-sum-exp.
   M-step: (M1) $\pi_r, \mathbf{\Sigma}_r$ (25); (M2) per-cluster wage WLS with $[V_{ir}]_{jj}$
   variance correction (26); (M3) measurement $(\boldsymbol{\mu}_r, \mathbf{\Sigma}_\nu)$ (27).
3. **Output**: $(\widehat{\mathbf{\Theta}}_1, \widehat{\mathbf{\Theta}}_2)$ and $\{\widehat{\omega}_{ir}, \widehat{\boldsymbol{m}}_{ir}, \widehat{V}_{ir}\}$.

Stage 2 — SMM on $\mathbf{\Theta}_3 = (\boldsymbol{b}_r, h_r)_{r=1}^R$ (Section 4.4.2).

1. **Freeze** Stage-1 outputs; compute the data mobility moments $\widehat{\boldsymbol{M}}$, type-weighted by $\widehat{\omega}_{ir}$.
2. **Repeat** until convergence of the criterion (16):
   Inner loop (per type, in parallel): solve the max Bellman fixed point (17),
   continuation by 1-D Gauss–Hermite quadrature (28);
   Simulation: simulate choices over careers using the argmax policy (11), wages under the objective law,
   beliefs by the subjective Kalman filter, common random numbers; form $\boldsymbol{M}(\mathbf{\Theta}_3)$;
   Outer step: update $(\boldsymbol{b}_r, h_r)$.
3. **Output**: $\widehat{\mathbf{\Theta}}_3$.

Return $(\widehat{\mathbf{\Theta}}_1, \widehat{\mathbf{\Theta}}_2, \widehat{\mathbf{\Theta}}_3, \widehat{\omega}_{ir})$.
