Mathematical Social Sciences 144 (2026) 102583

Contents lists available at ScienceDirect

Mathematical Social Sciences

journal homepage: www.elsevier.com/locate/mss

Matching with learning: a Bayesian assignment game approach

Saeed Najafi Zanganeh [a,*], Olivier Gossner [a,b]

[a] CREST, École Polytechnique, 91120, Palaiseau, France
[b] Department of Mathematics, London School of Economics, London, WC2A 2AE, UK

ARTICLE INFO

Keywords:
Two-sided matching
Bayesian learning
Assignment games
Core stability
Public signals

ABSTRACT

We study two-sided matching with transferable utility in which each match generates a surplus composed of a commonly known observable component and an initially unknown idiosyncratic component. The idiosyncratic component is revealed gradually through a sequence of publicly observed signals, one per worker firm pair per period. The paper examines a benchmark in which signals are produced externally and are observed regardless of the realized matching. This benchmark is motivated by markets with verifiable and portable evaluation such as standardized tests, certifications, platform ratings, and publication records. Under risk neutrality and common posteriors, each date's market reduces to a standard Shapley-Shubik assignment game on posterior expected surpluses, and stability is defined as membership in the core of that game. We provide closed-form Normal-Normal Bayesian updating. Core nonemptiness in the assignment game guarantees core outcomes at every date, and we derive comparative statics in signal precision and in the strength of industry specific comparative advantage. The main long run result is that, under a unique efficient matching margin condition, the matching that maximizes posterior expected total surplus coincides with the efficient matching after a finite random time almost surely, and the probability of selecting an inefficient matching at date t admits an explicit exponential upper bound for all sufficiently large dates. Simulations illustrate how information accumulation generates early turnover and how industry specific comparative advantage affects worker payoff dispersion. This framework gives a clear starting point to study more complex situations where information comes from actual matches or is kept private.

1. Introduction

Labor markets are matching markets: workers and firms form bilateral relationships that generate surplus, which is then divided through wages. Canonical models of two-sided matching typically assume that match values are known when relationships are formed, both in nontransferable utility environments (Gale and Shapley, 1962) and in transferable utility environments where stability is captured by the core of an assignment game (Shapley and Shubik (1971); see also Roth and Sotomayor (1992)).

Subsequent work studies structural extensions of assignment games. Camiña (2006) generalizes the model to settings where one seller owns multiple objects and characterizes the lattice structure of the core, while Atay and Núñez (2019) show that classical links between the core and stable sets do not extend straightforwardly to three-sided markets.

Yet a central feature of many labor markets is that match quality is only partially observed ex-ante (Jovanovic, 1979). Workers and firms observe *context*, credentials, job descriptions, industry affiliation, location, firm reputation, and infer likely fit, but important components of productivity and match specific compatibility are revealed only through interviews, probation periods, and realized performance.

A large literature models this informational friction by allowing match quality to be revealed only after relationships form. In labor economics, classic employer learning and job matching models emphasize that early outcomes and separations reflect gradual inference about worker job fit (e.g., Jovanovic (1979); Altonji and Pierret (2001); Lange (2007)). In matching theory, incomplete information raises the question of how to define and sustain stability when agents' beliefs differ or evolve; Liu et al. (2014) develop a notion of stability under incomplete information, and Bikhchandani (2017) studies stability with one-sided private information. Closest in spirit, Chen and Hu (2020) analyze how learning and matching interact when information is updated through the matching process itself.

Forges (2004) embeds private information about match dependent states into the assignment game and proves non-emptiness of the ex-ante and interim incentive compatible core, linking Bayesian incentive

* Corresponding author.
E-mail address: saeed.najafi@ip-paris.fr (S. Najafi Zanganeh).

https://doi.org/10.1016/j.mathsocsci.2026.102583
Received 26 February 2026; Received in revised form 16 August 2026; Accepted 7 September 2026
Available online 8 September 2026

compatibility to core stability in transferable utility markets. Doval and Schenone (2024) extend stability analysis to dynamic one to one matching markets, showing that consistent conjectures about continuation outcomes are sufficient for existence of stable allocations. Qi (2025) studies belief updating with endogenous public and private signals, characterizing asymptotic learning efficiency; although outside matching, such results clarify how information precision evolves over time.

1.1. Overview

This paper contributes a tractable benchmark that connects these strands: a transferable utility matching market with an observable component of surplus and Bayesian learning about an idiosyncratic component, in which information arrives publicly so that beliefs are common. The key simplification is as follows. At each date, the public signals received up to that date induce posterior means for the idiosyncratic components, and these posteriors define a valuation matrix. Because information is common, each date's market is a static Shapley-Shubik assignment game on that valuation matrix. The dynamics therefore live entirely in the evolution of valuations; the solution concept is unchanged across dates. This decoupling, a standard static solution concept at each date, driven by an exogenously updated valuation matrix, is what makes sharp long run results possible.

Our model integrates three ingredients. First, the observable component of surplus has an industry structure: firms belong to observable industries, and each worker has an industry specific comparative advantage. This generates correlation in observable surplus across firms in the same industry directly, through an additive term, rather than through a discrete choice specification of the kind used in the empirical matching literature following Choo and Siow (2006) and Galichon and Salanié (2022). Second, Bayesian learning about idiosyncratic quality: for each worker firm pair, an initially unknown idiosyncratic component is learned from an i.i.d. sequence of publicly observed signals; under Normal priors and Normal noise, posterior beliefs are available in closed form, yielding a tractable stochastic process for the valuation matrix. Third, stability with transfers: with transferable utility and public information, stability is defined as membership in the core of the assignment game defined by posterior expected surpluses, guaranteeing existence of stable outcomes at every date.

The paper delivers two sets of results. Statics: at each date, core outcomes exist; the posterior weight on realized signals increases with signal precision; and industry complementarities scale cross-industry differences in the observable component of surplus while leaving within-industry differences unaffected. Dynamics: under a mild *margin* condition guaranteeing that the efficient matching in the true surplus matrix is unique, the matching that maximizes posterior expected total surplus equals the efficient matching after some finite time almost surely. Moreover, we provide an explicit exponential upper bound, conditional on the margin, on the probability that the posterior optimal matching differs from the efficient one at sufficiently large dates. This formalizes the idea that information driven mismatches and turnover are transient when performance measures are informative and broadly observable.

The framework is intended as a benchmark. The public signal structure abstracts from adverse selection and strategic disclosure that arise under private learning, but it allows clean and transparent theory for large assignment problems with learning and transfers. The model is motivated by environments with verifiable and portable performance metrics, such as sales outcomes, standardized evaluations, publication records and platform ratings, where both sides and outside parties can observe comparable information and update beliefs in a disciplined way. Other settings with similar external scoring technology include academic job markets (publication record before revelation of departmental fit), professional sports scouting (combine metrics before on field performance), online labor platforms (verifiable credentials before task specific performance), and medical residency matching (board scores and evaluations before on the job fit).

A key implication of this information structure is that learning is exogenous to the matching path: the public signal process generates information about any potential pair (i,j) regardless of whether that pair is currently matched. This decoupling is deliberate. It provides a clear benchmark where equilibrium turnover and wage dynamics come from belief updating under common posteriors, while setting aside the more complex case where information mainly comes from realized matches. We discuss this matched pairs learning case and why it requires additional modeling ingredients in Section 8.4.

1.2. Roadmap

Section 2 presents the model environment. Section 3 specifies the observable component of surplus with industry structure. Section 4 formalizes Bayesian learning. Section 5 defines stability with transfers. Section 6 presents theoretical results. Section 7 provides simulation analysis. Section 8 discusses extensions and limitations. Section 9 concludes.

2. Model

2.1. Agents, matchings, and timing

The economy consists of a finite set of workers $I = \{1, \dots, n\}$ and a finite set of firms $J = \{1, \dots, m\}$. Each firm has capacity one and can employ at most one worker. Time is discrete, $t = 0, 1, 2, \dots$.

A (one-to-one) matching is a function $\mu : I \to J \cup \{\varnothing\}$ such that $|\mu^{-1}(j)| \leq 1$ for all $j \in J$, where $\mu(i) = \varnothing$ denotes unemployment. We also write $\mu(j) = i$ if firm j is matched to worker i, and $\mu(j) = \varnothing$ otherwise.

At each date t, the sequence of events is:

1. Information. Agents observe the public information set $\mathscr{F}_t$ (defined in Section 2.3), which includes all context variables and the history of public signals up to t.
2. Beliefs and valuations. Given $\mathscr{F}_t$, each worker firm pair (i,j) has a posterior expected surplus $v_{ij}(t) = \mathbb{E}[\pi_{ij} | \mathscr{F}_t]$.
3. Allocation and transfers. A matching μ_t and transfers (wages) are determined. We focus on outcomes that are stable (in the core) with respect to the valuation matrix $V(t) = (v_{ij}(t))$.
4. Signal arrival. At the end of the period (or equivalently, at the beginning of $t+1$), new public signals are realized, expanding $\mathscr{F}_{t+1}$.

This timing makes explicit that at date t the allocation and wages are formed conditional on the information available at that date.

2.2. Surplus and transferable utility

A match between worker i and firm j generates gross surplus

$$\pi_{ij} = u_{ij} + \varepsilon_{ij},$$

where u_{ij} is the observable component of the surplus, commonly known at date 0, and ε_{ij} is an idiosyncratic match quality component that is initially unknown. We refer to u_{ij} as the observable component throughout the paper.

We assume transferable utility: a wage w_{ij} transfers utility from the firm to the worker. Under risk neutrality, the relevant object for bargaining and stability at date t is the posterior expected surplus

$$v_{ij}(t) := \mathbb{E}[\pi_{ij} | \mathscr{F}_t] = u_{ij} + \mathbb{E}[\varepsilon_{ij} | \mathscr{F}_t].$$

We interpret wages as being negotiated at date t given the common public information $\mathscr{F}_t$. This can be justified in either of the following equivalent ways: (i) agents' contract on expected surplus and absorb realized fluctuations under risk neutrality; or (ii) contracts can be contingent on publicly verifiable information so that transfers effectively depend on $\mathscr{F}_t$. Under both interpretations, $v_{ij}(t)$ is the relevant valuation for the assignment problem at date t.

Let $U_i(t)$ denote worker i's payoff and $\Pi_j(t)$ denote firm j's payoff at date t. If i is matched with $j = \mu_t(i)$, we normalize the worker's payoff to the wage and the firm's payoff to residual surplus:

$$U_i(t) = w_{i,\mu_t(i)}, \Pi_{\mu_t(i)}(t) = v_{i,\mu_t(i)}(t) - w_{i,\mu_t(i)}.$$

Unmatched agents receive zero:

$$U_i(t) = 0 \text{ if } \mu_t(i) = \varnothing, \Pi_j(t) = 0 \text{ if } \mu_t(j) = \varnothing.$$

Equivalently, for any matched pair (i,j) at date t,

$$U_i(t) + \Pi_j(t) = v_{ij}(t),$$

which is the feasibility condition used in the assignment game.

2.3. Information and beliefs

Let $\mathscr{F}_t$ denote the public information available at date t. It contains:

1. Context: all observable variables that determine. $\{u_{ij}\}_{(i,j)\in I\times J}$
2. Signal history: the collection of public signals $\{s_{ij}(\tau)\}_{\tau=1}^{t}$ for all $(i,j) \in I \times J$.

Formally, for each pair (i,j) and each date $t \geq 1$, the public signal takes the form

$$s_{ij}(t) = \varepsilon_{ij} + \eta_{ij}(t),$$

where ε_{ij} is the idiosyncratic component of the surplus $\pi_{ij} = u_{ij} + \varepsilon_{ij}$ introduced in Section 2.2, and $\eta_{ij}(t)$ is zero mean noise, independent across (i,j,t) and independent of ε_{ij}. The prior on ε_{ij} and the distribution of $\eta_{ij}(t)$ are specified in Section 4, where the posterior update is derived in closed form. At this stage, the only feature we use is that $\mathscr{F}_t$ determines common posterior means $\mu_{ij,t} = \mathbb{E}[\varepsilon_{ij}|\mathscr{F}_t]$ for every pair.

Signals are publicly observed: for each pair (i,j) and each date $t \geq 1$, the realization $s_{ij}(t)$ is observed by both sides and by all market participants, so posteriors are common and the date-t allocation problem remains a complete information assignment game. In the baseline benchmark, signals are observed for all worker firm pairs at all dates, independently of whether i and j are matched at date t. Equivalently, $\mathscr{F}_t$ is generated by context and the full signal history and does not depend on the realized matching sequence $\{\mu_\tau\}_{\tau\leq t}$. Signals arrive exogenously, and we do not model endogenous information acquisition, experimentation, or strategic disclosure. This information structure is strong. Evaluation systems that a whole market can observe usually produce measures at the worker level, or they record performance inside matches that were actually formed. A signal attached to a pair that never met is a more demanding object than such systems normally deliver. We adopt it because it keeps each date inside a standard assignment game and makes the learning process transparent, and not because we regard it as a description of any particular evaluation system. In many applications, signals would be observed only for interviewed pairs or realized matches, and Section 8.4 discusses how such restricted observability couples learning and matching.

3. The observable component of surplus

This section specifies the observable component of match surplus, u_{ij}. The goal is to allow for correlated observable valuations among firms that belong to the same industry, capturing the idea that a worker who fits one firm in an industry tends to fit other firms in the same industry as well.

3.1. Industry partition

Firms are partitioned into a finite set of industries $= \{1,\dots,\bar{G}\}$. Each firm $j \in J$ belongs to exactly one industry, denoted $g(j) \in G$. Industries can be interpreted as sectors (e.g., technology, finance, manufacturing) or occupational clusters (e.g., engineering, sales, management).

Workers may have industry specific comparative advantage: holding constant firm level characteristics, a worker can be systematically better suited to one industry than another. We model this through worker industry interaction terms that load common components across all firms within the same industry.

3.2. Parametric specification with industry structure

We now specialize the observable component of surplus u_{ij} to a parametric form with industry structure, which we use for the comparative statics in Section 6.2 and for the simulations in Section 7. We emphasize that the general theoretical results of Sections 5 and 6.1 and the convergence result Proposition 4 require only that u_{ij} is commonly known at date 0, and do not depend on any parametric form. The specification below is introduced as a concrete example that is convenient for interpretation and simulation.

Let

- $z_{ij} \in \mathbb{R}^{d_z}$ be match level covariates (e.g., skill requirement fit, distance, occupation match),
- $x_j \in \mathbb{R}^{d_x}$ be observed firm characteristics (e.g., size, prestige, technology),
- $v_i \in \mathbb{R}^{d_v}$ be observed worker characteristics (e.g., education, experience, credentials).

Define the baseline observable index

$$V_{ij} := z_{ij}'\beta + x_j'\delta + v_i'\xi,$$

where β, δ, ξ are parameter vectors of conformable dimensions.

To incorporate industry structure, let $\theta_{ig} \in \mathbb{R}$ denote worker i's comparative advantage in industry g, normalized such that $\mathbb{E}[\theta_{ig}] = 0$ (across workers within each industry). The observable component of surplus is

$$u_{ij} = V_{ij} + \lambda\,\theta_{i,g(j)} + \alpha_j + \gamma_i,$$

where $\lambda \geq 0$ measures the strength of industry specific comparative advantage, α_j is a firm fixed effect, and γ_i is a worker fixed effect.

The additive fixed effects (α_j, γ_i) capture time invariant firm and worker heterogeneity in observable component of surplus. The term $\lambda\theta_{i,g(j)}$ generates correlation across all firms in the same industry from the perspective of a given worker.

3.3. Implications for substitutability and sorting

The specification yields transparent comparative statics in λ:

- If $\lambda = 0$, industry membership plays no role beyond what is already captured in observed covariates and fixed effects; observable component of surplus varies only through $(z_{ij}, x_j, v_i, \alpha_j, \gamma_i)$.
- If $\lambda > 0$, industry specific worker components enter observable surplus and can affect cross industry sorting when they are large enough relative to the remaining surplus terms and assignment constraints. For fixed i, the term $\theta_{i,g(j)}$ is common across all firms in the same industry. It therefore cancels from within industry surplus differences and affects comparisons across industries.

This reduced form structure is deliberately streamlined. It captures correlation in observable surplus within each industry through an additive term rather than through a discrete choice micro foundation. The parameters $(\beta,\delta,\xi,\lambda)$, together with industry labels $g(j)$, are of a form that is compatible with cross-sectional variation used in the empirical matching literature on transferable utility models; we discuss empirical relevance further in Section 8.2, while stressing that the paper is theoretical and does not formulate an identification result.

4. Bayesian learning about latent match quality

This section specifies how agents learn about ε_{ij} over time from public signals and derives the posterior expected surplus $v_{ij}(t)$ that enters the date t assignment game.

4.1. Prior

For each pair (i,j), latent match quality ε_{ij} is initially unknown. We assume independent priors across pairs with common Normal prior:

$\varepsilon_{ij} \sim \mathcal{N}(m_0, \tau_0^2)$, where $m_0 \in \mathbb{R}$ and $\tau_0^2 > 0$ are common knowledge.

The baseline information structure combines two distinct assumptions. The prior on $(\varepsilon_{ij})_{(i,j)\in I\times J}$ is independent across pairs, and the signal $s_{ij}(t)$ is publicly observed for every pair at every date, regardless of the realized matching. The two assumptions can be relaxed separately. Section 4.6 presents a component extension $\varepsilon_{ij} = a_i + b_j + \zeta_{ij}$ under which almost sure eventual selection, Proposition 4(a), continues to hold, even though the separate components are not identified. The explicit exponential constant in Proposition 4(b) uses independence. Public observability is the defining feature of the benchmark and is maintained throughout; the matched pair learning alternative in which information is endogenous to the matching sequence is discussed in Section 8.4.

4.2. Public signal technology

We specify the noise in the public signal introduced in Section 2.3 as Gaussian:

$$\eta_{ij}(t) \sim \mathcal{N}(0, \sigma^2),$$

independent across (i,j,t) and independent of ε_{ij}. The parameter $\sigma^2 > 0$ governs the precision of the signal; smaller σ^2 corresponds to more precise information. Together with the Gaussian prior of Section 4.1, this specification delivers closed form posterior updating, developed in Section 4.3.

4.3. Posterior updating

Lemma 1. ***(Normal-Normal conjugacy).*** *After observing the history $\{s_{ij}(\tau)\}_{\tau=1}^{t}$ for pair (i,j), the posterior distribution is Normal:*

$$\varepsilon_{ij} | \mathcal{F}_t \sim \mathcal{N}\left(\mu_{ij,t}, \tau_t^2\right),$$

with posterior mean

$$\mu_{ij,t} = \frac{\tau_0^{-2} m_0 + \sigma^{-2} \sum_{\tau=1}^{t} s_{ij}(\tau)}{\tau_0^{-2} + t\sigma^{-2}} = (1-\kappa_t) m_0 + \kappa_t \bar{s}_{ij,t},$$

and posterior precision

$$\tau_t^{-2} = \tau_0^{-2} + t\sigma^{-2},$$

where $\bar{s}_{ij,t} := \frac{1}{t}\sum_{\tau=1}^{t} s_{ij}(\tau)$ is the sample mean and

$$\kappa_t := \frac{t\tau_0^2}{t\tau_0^2 + \sigma^2} \in (0,1)$$

is the (Kalman) weight on the sample mean.

The posterior variance τ_t^2 depends only on t (common across pairs under the homoscedastic signal structure), while $\mu_{ij,t}$ is pair specific through $\bar{s}_{ij,t}$.

4.4. Learning properties

Proposition 1. ***(Learning dynamics).***

(a) (Consistency) $\mu_{ij,t} \to \varepsilon_{ij}$ almost surely as $t \to \infty$.
(b) (Uncertainty resolution) $\tau_t^2 = \left(\tau_0^{-2} + t\sigma^{-2}\right)^{-1} = O(1/t)$.
(c) (Precision comparative statics) Let $p := \sigma^{-2}$ denote signal precision. For every positive date, $\kappa_t = \frac{t\tau_0^2 p}{t\tau_0^2 p+1}$ is strictly increasing in p.

Proof.

(a) Since $\mathbb{E}[s_{ij}(\tau)|\varepsilon_{ij}] = \varepsilon_{ij}$, the strong law implies $\bar{s}_{ij,t} \to \varepsilon_{ij}$ a.s. Also $\kappa_t \to 1$. Hence $\mu_{ij,t} = (1-\kappa_t) m_0 + \kappa_t \bar{s}_{ij,t} \to \varepsilon_{ij}$ a.s.
(b) Immediate from $\tau_t^{-2} = \tau_0^{-2} + t\sigma^{-2}$.
(c) With $p = \sigma^{-2}$, $\kappa_t = \frac{t\tau_0^2 p}{t\tau_0^2 p+1}$ and $\frac{\partial \kappa_t}{\partial p} = \frac{t\tau_0^2}{(t\tau_0^2 p+1)^2} > 0$.

4.5. Posterior expected surplus and the date-t assignment game

Given the decomposition $\pi_{ij} = u_{ij} + \varepsilon_{ij}$, the posterior expected surplus at date t is

$$v_{ij}(t) := \mathbb{E}[\pi_{ij} | \mathcal{F}_t] = u_{ij} + \mu_{ij,t}.$$

Let $V(t) = (v_{ij}(t))_{i\in I, j\in J}$ denote the valuation matrix. At each date t, the market therefore induces a standard transferable utility assignment game with valuations $V(t)$. As signals accumulate, the matrix $V(t)$ evolves, potentially altering the identity of stable matchings and equilibrium wages.

Key feature. Under Normal-Normal learning, $\mu_{ij,t}$ is linear in the realized signals (through $\bar{s}_{ij,t}$), which yields clean comparative statics and facilitates the convergence analysis in later sections.

4.6. Extension remark: correlated latent match quality

The baseline assumes ε_{ij} are independent across pairs, which yields scalar updating and a common posterior variance τ_t^2 across (i,j). This independence is not essential for the main convergence logic. For instance, one may allow latent quality to share worker and firm level components, e.g.

$$\varepsilon_{ij} = a_i + b_j + \zeta_{ij}$$

where a_i captures an initially unknown worker level component, b_j an initially unknown firm level component, and ζ_{ij} a residual pair specific component, with a Normal joint prior and Normal signal noise. Throughout, ζ_{ij} denotes a latent residual and is distinct from the parameter vector ξ of Section 3.2, which remains a coefficient on worker covariates.

Under the same public signal benchmark, with signals observed for all (i,j) at all dates, the joint prior over (a_i, b_j, ζ_{ij}) is Gaussian and the posterior remains Gaussian after observing the signal history. Repeated signals for a given pair reveal the sum ε_{ij} rather than its separate parts, so the individual components need not be identified from the signal history alone. This is not needed for our purposes. Proposition 4 uses only the convergence of the posterior mean of ε_{ij}, and that posterior mean still converges almost surely to its true value. The posterior expected surplus matrix therefore converges entrywise to the true surplus matrix, and

Proposition 4(a) continues to hold.

5. Stability and wages in the assignment game

At each date t, posterior expected surpluses $v_{ij}(t)$ define a transferable utility assignment game. This section states the equilibrium concept (the core), records existence, and clarifies how we compute core selections in the simulations.

5.1. Outcomes and feasibility

Fix a date t and the valuation matrix $(t) = \left(v_{ij}(t)\right)_{i\in I, j\in J}$. An outcome is a triple (μ, U, Π), where μ is a matching and $U = (U_i)_{i\in I}$, $\Pi = \left(\Pi_j\right)_{j\in J}$ are payoff vectors for workers and firms.

An outcome is feasible (at date t) if there exist transfers (wages) such that total expected surplus is allocated within each matched pair and unmatched agents obtain their outside option normalized to zero. Equivalently, feasibility requires:

$$U_i + \Pi_{\mu(i)} = v_{i,\mu(i)}(t) \text{ for each matched } i,$$

and

$$U_i = 0 \text{ if } \mu(i) = \varnothing, \Pi_j = 0 \text{ if } \mu^{-1}(j) = \varnothing.$$

Given a feasible outcome, the associated wage paid by firm $j = \mu(i)$ to worker i can be taken as

$$w_{ij} = U_i,$$

so that the firm's payoff is $\Pi_j = v_{ij}(t) - w_{ij}$.

5.2. Blocking and myopic core stability

A worker firm pair (i, j) blocks a feasible outcome (μ, U, Π) at date t if they can form a match and choose a transfer w that makes both strictly better off. Under transferable utility, this is equivalent to:

$$U_i + \Pi_j < v_{ij}(t).$$

Indeed, when the strict inequality holds, there exists $w \in \left(U_i, v_{ij}(t) - \Pi_j\right)$ such that worker i obtains $w > U_i$ and firm j obtains $v_{ij}(t) - w > \Pi_j$.

Definition 1. (Myopic Core Stability). An outcome (μ, U, Π) is myopically core stable at date t if it is feasible and satisfies the no-blocking inequalities

$$U_i + \Pi_j \geq v_{ij}(t) \ \forall (i, j) \in I \times J.$$

The set of all myopic core stable outcomes at date t is the core of the assignment game induced by $V(t)$.

Lemma 2. ***(Core nonemptiness; Shapley-Shubik).*** *For any finite valuation matrix $V(t)$, the core of the assignment game is nonempty.*

Proof.

Consider the primal linear program that maximizes total surplus over feasible matchings (the max weight bipartite matching problem) and its dual that selects (U, Π) to minimize $\sum_i U_i + \sum_j \Pi_j$ subject to $U_i + \Pi_j \geq v_{ij}(t)$ for all (i, j). Strong duality implies equality of optimal values. Complementary slackness yields an optimal matching μ_t and dual payoffs (U, Π) satisfying feasibility on matched pairs and no-blocking on all pairs. The dual payoff variables are restricted to be nonnegative, reflecting the zero outside option.

5.3. Core selection and computation

The core is typically a lattice; comparative statics and simulations therefore require a selection rule. We focus on the worker optimal core outcome (analogously, one could compute the firm optimal outcome).

Operationally, at each date t:

1. Efficient matching. Compute an efficient matching μ_t that maximizes total expected surplus:

$$\mu_t \in \operatorname*{argmax}_{\mu} \sum_{i:\mu(i)\neq\varnothing} v_{i,\mu(i)}(t).$$

2. Worker optimal payoffs in the core. Among core payoffs consistent with μ_t, select the worker optimal one by solving

$$\max_{U,\Pi} \sum_{i\in I} U_i$$

subject to

$$U_i + \Pi_j \geq v_{ij}(t) \forall (i, j) \in I \times J,$$

$$U_i + \Pi_{\mu_t(i)} = v_{i,\mu_t(i)}(t) \qquad \forall i \in I \textit{ with } \mu_t(i) \neq \varnothing,$$

$$U_i \geq 0, \ \Pi_j \geq 0 \forall i, j.$$

Recall from Section 5.1 that an unmatched worker receives $U_i = 0$ and an unmatched firm receives $\Pi_j = 0$. These zero payoff equalities are imposed as constraints for every unmatched worker and firm. The nonnegativity constraints $U_i \geq 0$ and $\Pi_j \geq 0$ therefore impose individual rationality relative to that outside option. This selection is standard and yields a unique payoff vector under generic conditions.

6. Comparative statics: precision and industry structure

This section records two transparent comparative statics that follow directly from the Normal-Normal learning structure and from the industry specification of observable component of surplus. These results will be used to interpret the simulation patterns and the convergence analysis.

6.1. Signal precision and responsiveness of posterior values

Let $p := \sigma^{-2}$ denote signal precision and recall $\kappa_t = \frac{t\tau_0^2 p}{t\tau_0^2 p + 1}$.

Proposition 2. *(Precision increases responsiveness).*

(a) κ_t is strictly increasing in p.

(b) Fix t and two signal histories for (i, j) with distinct sample means $\bar{s}_{ij,t} \neq \bar{s}'_{ij,t}$. Then

$$\mu_{ij,t}(\bar{s}') - \mu_{ij,t}(\bar{s}) = \kappa_t(\bar{s}' - \bar{s}),$$

so $\left|\mu_{ij,t}(\bar{s}') - \mu_{ij,t}(\bar{s})\right|$ is strictly increasing in p.

Proof.

(a) Direct differentiation gives $\frac{\partial\kappa_t}{\partial p} = \frac{t\tau_0^2}{\left(t\tau_0^2 p + 1\right)^2} > 0$.

(b) Since $\mu_{ij,t} = (1 - \kappa_t)m_0 + \kappa_t \bar{s}_{ij,t}$, differences scale linearly with κ_t, which is increasing in p by part (a).

Corollary 1. *Higher precision (smaller σ^2) makes it more likely that learning overturns purely context driven rankings: for fixed u_{ij}, sufficiently favorable signals can raise $v_{ij}(t) = u_{ij} + \mu_{ij,t}$ above alternatives with higher u but unfavorable signals.*

6.2. Industry complementarities and cross industry differentiation

Recall

$$u_{ij} = V_{ij} + \lambda\theta_{i,g(j)} + \alpha_j + \gamma_i.$$

Proposition 3. (Industry structure).

Fix worker i. For two firms j, k in the same industry g,

$$u_{ij} - u_{ik} = \left(V_{ij} - V_{ik}\right) + \left(\alpha_j - \alpha_k\right),$$

which is independent of λ. For two firms $j \in g$ and $j' \in g' \neq g$,

$$u_{ij} - u_{ij'} = \left(V_{ij} - V_{ij'}\right) + \lambda\left(\theta_{ig} - \theta_{ig'}\right) + \left(\alpha_j - \alpha_{j'}\right),$$

which varies linearly with λ.

Proof.

Immediate by differencing u_{ij} across firms and noting that γ_i cancels. Within an industry, $\theta_{i,g(j)} = \theta_{i,g(k)}$, so the λ-term cancels; across industries it does not.

Proposition 3 is a statement about differences in the observable component of surplus, not about equilibrium wages. Mapping surplus differences into worker payoff dispersion depends on the choice of core selection and on the full distribution of primitives; we do not pursue such a mapping analytically. Section 7.4 reports the corresponding numerical payoff decomposition under the worker optimal core selection. This decomposition is a numerical outcome and does not follow from Proposition 3.

6.3. Convergence and stability

Let the true surplus matrix be

$$W_{ij} = u_{ij} + \varepsilon_{ij},$$

and the posterior expected surplus matrix at date t be

$$\widehat{W}_{ij,t} = u_{ij} + \mu_{ij,t},$$

where $\mu_{ij,t} = \mathbb{E}\left[\varepsilon_{ij}\middle|\mathscr{F}_t\right]$ under Normal-Normal updating.

For any matching μ, define the total surplus

$$S(W,\mu) := \sum_{i\in I:\,\mu(i)\neq\varnothing} W_{i,\mu(i)}.$$

Let μ^* be an efficient matching for W, and define the margin

$$\Delta := \min_{\mu\neq\mu^*} \{S(W,\mu^*) - S(W,\mu)\}.$$

We assume $\Delta > 0$, i.e., the efficient matching is unique with positive separation. We treat Δ as a fixed constant. Note that in large markets ($|I|$, $|J|\to\infty$), the minimum gap between the efficient and second best matching may shrink. If $\Delta\to 0$, the convergence time bounds derived below would naturally increase.

Let μ_t denote a matching that maximizes posterior expected total surplus:

$$\mu_t \in \operatorname*{argmax}_{\mu} S(\widehat{W}_t,\mu).$$

Proposition 4. ***(Eventual Selection of the Efficient Matching and Exponential Error Bound)***

Fix a realization of ε such that the true surplus matrix W has a unique efficient matching μ^* with margin $\Delta > 0$. Then:

(a) Almost sure eventual selection. With probability one over the signal process $\left\{s_{ij}(\tau)\right\}$ conditional on ε, there exists a finite random time $T < \infty$ such that $\mu_t = \mu^*$ for all $t \geq T$.

(b) Explicit exponential tail bound. There exists a finite t_0, depending on ε, Δ, ${t_0}^2$ and σ^2, such that for every $t \geq t_0$,

$$\mathbb{P}(\mu_t \neq \mu^* \mid \varepsilon) \leq 2\left|I\right|\left|J\right| \exp\left(-\frac{\Delta^2}{32\,|I|^2\,\sigma^2}\,t\right),$$

where σ^2 is the signal noise variance in $s_{ij}(t) = \varepsilon_{ij} + \eta_{ij}(t)$.

Proof.

Define the estimation error matrix E_t with entries

$$E_{ij,t} := \widehat{W}_{ij,t} - W_{ij} = \mu_{ij,t} - \varepsilon_{ij}.$$

Step 1. (stability under uniform perturbations).

For any matching μ,

$$S(\widehat{W}_t,\mu) - S(W,\mu) = \sum_{(i,j)\in\mu} E_{ij,t}.$$

Since any matching uses at most $|I|$ edges, we have

$$|S(\widehat{W}_t,\mu) - S(W,\mu)| \leq |I| \quad \|E_t\|_\infty,$$

where $\| E_t\|_\infty = \max_{i,j}|E_{ij,t}|$.

Hence, for any $\mu \neq \mu^*$,

$$S(\widehat{W}_t,\mu^*) - S(\widehat{W}_t,\mu) \geq \underbrace{S(W,\mu^*) - S(W,\mu)}_{\geq\Delta} - 2|I| \quad \|E_t\|_\infty.$$

Therefore, if $\| E_t\|_\infty < \frac{\Delta}{(2|I|\ )}$, then

$$S(\widehat{W}_t,\mu^*) - S(\widehat{W}_t,\mu) > 0 \qquad \forall\mu \neq \mu^*,$$

so μ^* is the unique maximizer under $\widehat{W}_t$ and thus $\mu_t = \mu^*$.

Equivalently,

$$\{\mu_t \neq \mu^*\} \subseteq \left\{\ \| E_t\|_\infty \geq \frac{\Delta}{(2|I|\ )}\right\} \tag{1}$$

Step 2. (distribution of $E_{ij,t}$).

Under Normal-Normal learning with i.i.d. signal noise $\eta_{ij}(t) \sim N\big(0, \sigma^2\big)$, we have

$$\mu_{ij,t} = (1-\kappa_t)m_0 + \kappa_t\bar{s}_{ij,t}, \quad \bar{s}_{ij,t} = \frac{1}{t}\sum_{\tau=1}^{t} s_{ij}(\tau), \quad \kappa_t = \frac{t\tau_0^2}{t\tau_0^2 + \sigma^2} \in (0,1).$$

Fix a pair (i,j). Conditional on ε, the posterior error admits the decomposition

$$E_{ij,t} := \mu_{ij,t} - \varepsilon_{ij} = (1-\kappa_t)\big(m_0 - \varepsilon_{ij}\big) + \kappa_t\left(\bar{s}_{ij,t} - \varepsilon_{ij}\right).$$

The second term is Gaussian conditional on ε. Indeed,

$$\bar{s}_{ij,t} - \varepsilon_{ij} = \frac{1}{t}\sum_{\tau=1}^{t}\eta_{ij}(\tau) \sim N\left(0,\frac{\sigma^2}{t}\right),$$

so

$$\kappa_t\left(\bar{s}_{ij,t} - \varepsilon_{ij}\right)\Big|\varepsilon \sim N\left(0,\kappa_t^2\frac{\sigma^2}{t}\right) \text{and hence } \operatorname{Var}\left(\kappa_t\left(\bar{s}_{ij,t} - \varepsilon_{ij}\right)\Big|\varepsilon\right) \leq \frac{\sigma^2}{t}.$$

The first term is a (non-random) bias conditional on ε. Let

$$B(\varepsilon) := \max_{(i,j)\in I\times J}\left|m_0 - \varepsilon_{ij}\right| < \infty.$$

Then, for all (i,j),

$\left|(1-\kappa_t)\big(m_0 - \varepsilon_{ij}\big)\right| \leq (1-\kappa_t)\,B(\varepsilon)$.

Fix the threshold $x := \frac{\Delta}{(2|I|\ )}$. Since $1 - \kappa t = \frac{\sigma^2}{t\tau_0^2 + \sigma^2}$ tends to zero as t grows, there exists a finite t_0 such that $(1-\kappa_t)B(\varepsilon) \leq \frac{x}{2}$ for all $t \geq t_0$. For such t, the triangle inequality implies

$$\mathbb{P}\left(\left|E_{ij,t}\right| \geq x \mid \varepsilon\right) \leq \mathbb{P}\left(\left|\kappa_t\left(\bar{s}_{ij,t} - \varepsilon_{ij}\right)\right| \geq x/2 \mid \varepsilon\right).$$

Applying the standard Gaussian tail bound and using $\mathrm{Var}(\cdot|\varepsilon) \leq \frac{\sigma^2}{t}$ yields

$$\mathbb{P}\left(\left|E_{ij,t}\right| \geq x \mid \varepsilon\right) \leq 2\exp\left(-\frac{x^2 t}{8\sigma^2}\right) \quad (2)$$

Step 3. (union bound over pairs).

By a union bound over all $|I||J|$ pairs,

$$\mathbb{P}(\| E_t\|_\infty \geq x \mid \varepsilon) \leq \sum_{i\in I}\sum_{j\in J}\mathbb{P}\left(\left|E_{ij,t}\right| \geq x \mid \varepsilon\right) \leq 2\left|I\right|\left|J\right|\exp\left(-\frac{x^2 t}{8\sigma^2}\right) \quad (3)$$

Set $x = \frac{\Delta}{(2|I|\ )}$ and combine (1) and (3) to obtain

$$\mathbb{P}(\mu_t \neq \mu^* \mid \varepsilon) \leq 2\left|I\right|\left|J\right|\exp\left(-\frac{\Delta^2}{32|I|^2\sigma^2}t\right),$$

which proves part (b).

Step 4. (almost sure eventual selection).

The bound in part (b) holds for all $t \geq t_0$ and is summable over those dates. What happens at the finitely many dates before t_0 does not affect the argument, so the Borel-Cantelli lemma implies that $\{\mu_t\mu^*\}$ occurs only finitely often almost surely, conditional on ε. This yields part (a).

7. Simulation analysis

This section illustrates the model's implications in a controlled environment. The goal is not calibration, but to (i) visualize how Bayesian learning reshapes the valuation matrix $V(t)$, (ii) quantify induced turnover under repeated re-optimization of the assignment, and (iii) document how information precision and industry structure affect equilibrium payoffs under a fixed core selection (the worker optimal core).

7.1. Data generating process

We consider balanced markets with

$n = m = 100,$

and partition firms into $G = 5$ industries with 20 firms each. Each firm j belongs to a unique industry $g(j) \in \{1, \ldots, G\}$.

For each worker i and firm j, we draw:

- worker skill vectors $s_i \sim \mathcal{N}(0, I_3)$ and firm requirement vectors $r_j \sim \mathcal{N}(0, I_3)$;

(Here s_i denotes worker skill vectors, while $s_{ij}(t)$ denotes performance signals; the distinction is by indices.)

- a match-specific covariate

$$z_{ij} = -\frac{1}{2}\ \|s_i - r_j\|^2;$$

- firm covariates $x_j \sim \mathcal{N}(0,1)$ and worker covariates $v_i \sim \mathcal{N}(0,1)$;
- firm fixed effects $\alpha_j \sim \mathcal{N}(0, 0.3)$ and worker fixed effects $\gamma_i \sim \mathcal{N}(0, 0.3)$, with 0.3 denoting the standard deviation in both distributions;
- worker-industry comparative advantage shocks $\theta_{ig} \sim \mathcal{N}(0,1)$, independent across (i, g).

Given parameter values

$\beta = 0.5, \delta = 0.4, \xi = 0.3, \lambda = 0.5,$

the observable component is

$u_{ij} = \beta z_{ij} + \delta x_j + \xi v_i + \lambda\, \theta_{i,g(j)} + \alpha_j + \gamma_i.$

Latent match quality is drawn independently across pairs:

$\varepsilon_{ij} \sim \mathcal{N}(0,1).$

Agents share a common Normal prior $\mathcal{N}(m_0, \tau_0^2)$ with $m_0 = 0$ and $\tau_0^2 = 1$. At each date $t \geq 1$, a public signal is realized:

$s_{ij}(t) = \varepsilon_{ij} + \eta_{ij}(t), \eta_{ij}(t) \sim \mathcal{N}(0, \sigma^2),$

independent across (i, j, t). In the baseline we set $\sigma^2 = 0.4$.

After observing $\{s_{ij}(\tau)\}_{\tau=1}^{t}$, agents compute the posterior mean $\mu_{ij,t}$ (Section 4) and form posterior expected surplus

$v_{ij}(t) = u_{ij} + \mu_{ij,t}$

At each t, we:

1. At each date t, we allow agents to remain unmatched ($\varnothing$) with outside option normalized to zero. Computationally, we implement unemployment by augmenting the assignment problem with dummy firms (and, symmetrically, dummy workers) so that matching to a dummy partner yields value 0. We then solve the resulting maximum weight *perfect* matching problem (Hungarian algorithm) on the augmented matrix. The induced matching on the original worker firm sets is the efficient matching $\mu_t$for the economy with unemployment, i.e.,

$$\mu_t \in \operatorname*{argmax}_{\mu} \sum_{i:\mu(i)\neq\varnothing} v_{i,\mu(i)}(t),$$

where pairs assigned to dummy partners are recorded as $\mu_t(i) = \varnothing$

2. select payoffs using the worker optimal core associated with μ_t (Section 5.3), and record worker payoffs $(U_i)_{i'I}$ $(U_i)_{i\in I}$.

We simulate 50 independent markets and report across market averages. Random draws are generated with NumPy's legacy RandomState generator using seed 2024 for reproducibility.

7.2. Baseline dynamics

We report three summary statistics.

Let Turnover(t) denote the fraction of workers whose assignment changes between μ_{t-1} and μ_t, counting changes into or out of unemployment:

$\text{Turnover}(t) = \frac{1}{n}\left|\{i \in I : \mu_t(i) \neq \mu_{t-1}(i)\}\right|.$

The posterior standard deviation is $\tau_t = \sqrt{\tau_t^2}$, where $\tau_t^2 = \left(\tau_0^{-2} + t\sigma^{-2}\right)^{-1}$.

Let $\mu^{\text{ctx}} \in \operatorname{argmax}_{\mu}\sum_{i:\mu(i)\neq\varnothing} u_{i,\mu(i)}$, where unemployment is handled in the same way as above (dummy partners with value 0).

At $t = 0$, we set μ_0 to be the efficient matching under prior means, i.e., using $v_{ij}(0) = u_{ij} + m_0$ with $m_0 = 0$, and computing the corresponding worker optimal core payoffs. Table 1 reports the evolution of turnover, posterior uncertainty, and overlap with the context only assignment.

We decompose dispersion of worker payoffs U_i into within- and between-industry components using the industry of the worker's matched firm. Let $g_i = g(\mu_t(i))$ for employed workers and exclude unemployed workers from the industry decomposition. "Within-industry

Table 1
Learning and matching dynamics (baseline).

t	Turnover (%)	Posterior SD τ_t	Overlap with context-only (%)
1	81.2	0.535	18.8
5	23.1	0.272	15.8
10	12.3	0.196	15.2
20	7.2	0.140	15.1

Notes. Turnover is the percentage of workers whose assignment changes from $t-1$ to t. τ_t is computed from the conjugate Normal update. Overlap compares μ_t to the maximizer of $\sum_i u_{i,\mu(i)}$. Entries are averages over 50 markets. Employment rises from 62.0 percent at date 0 to 87.5 percent at date 1 and 90.3 percent at date 20, so first period turnover includes moves into and out of employment as well as job to job reallocation.

SD" is the standard deviation of $U_i - \mathbb{E}[U_i|g_i]$ across employed workers. "Cross-industry SD" is the employment weighted standard deviation of the industry means $\mathbb{E}[U_i|g_i]$. Table 2 summarizes the within- and cross-industry dispersion of worker payoffs over time.

Three patterns stand out. First, turnover is high in the first period and then declines as posterior uncertainty shrinks. Second, overlap with the context only assignment falls immediately after signals arrive and then remains close to 15 percent. Third, within industry payoff dispersion rises more than cross industry dispersion, reflecting that learning reveals idiosyncratic components of match quality that are not aligned with industry segmentation. These finite horizon simulations describe reallocation dynamics and are not a quantitative test of Proposition 4's asymptotic tail bound.

7.3. Sensitivity to signal precision

We vary signal noise

$\sigma^2 \in \{0.2, 0.4, 0.8, 1.6\}$,

Lower σ^2 corresponds to more precise signals and (by Proposition 2) stronger responsiveness of posterior means to realized performance.

Table 3 reports how signal precision affects first period turnover, context overlap at date 20, and within industry payoff dispersion at date 20.

Higher signal precision generates somewhat larger first period reallocation. Reallocation is high at every precision level in this design, so the effect is modest. We do not find a systematic effect of precision on the date 20 overlap or on within industry payoff dispersion.

7.4. Sensitivity to industry complementarities

We vary the strength of industry specific comparative advantage

$\lambda \in \{0, 0.25, 0.5, 1.0\}$,

holding $\sigma^2 = 0.4$ fixed. Recall from Proposition 3 that increasing λ magnifies cross-industry differences in u_{ij} while leaving within-industry differences in u_{ij} unaffected.

Table 4 reports how industry complementarity affects worker payoff dispersion at $t = 20$.

Table 2
Dispersion of worker payoffs Ui (baseline).

t	Within-industry SD	Cross-industry SD	Ratio (within/cross)
0	0.37	0.08	4.39
5	0.65	0.13	5.06
10	0.65	0.13	5.13
20	0.65	0.13	5.08

Notes. SDs are computed from worker optimal core payoffs. Industry is defined by the matched firm's industry. Averages over 50 markets. The reported ratio is computed from the unrounded across market mean SDs.

Table 3
Effect of Signal Precision.

	Turnover from $t = 0$ to 1 (%)	Overlap at $t = 20$ (%)	Within-industry SD at $t = 20$
0.2	83.2	14.7	0.65
0.4	81.2	15.1	0.65
0.8	78.0	15.2	0.65
1.6	73.3	15.5	0.64

Notes. "Turnover from $t = 0$ to 1" compares μ_0 to μ_1. "Final overlap" and "final within-industry SD" are computed at $t = 20$. Averages over 50 markets.

Table 4
Effect of Industry Complementarity (payoffs at t = 20).

λ	Within-industry SD	Cross-industry SD	Ratio (within/cross)
0	0.59	0.12	4.89
0.25	0.60	0.12	4.97
0.5	0.65	0.13	5.08
1.0	0.83	0.16	5.16

Notes. SDs computed from worker-optimal core payoffs at $t = 20$. Averages over 50 markets. The reported ratio is computed from the unrounded across market mean SDs.

As industry complementarity increases, both within industry and cross industry worker payoff dispersion rise, while their ratio changes little. Proposition 3 concerns observable surplus differences across firms for a fixed worker, whereas Table 4 measures equilibrium payoff dispersion across heterogeneous workers. Because industry specific comparative advantage varies across workers within the same industry, a larger complementarity parameter can raise both components.

8. Discussion

8.1. Economic interpretation and mechanisms

Matching stability changes as beliefs change. An allocation that is stable under current beliefs may become blocked after new information arrives. In the model, turnover therefore arises from belief updating rather than from exogenous shocks. Learning driven turnover is front loaded. It is highest when posterior uncertainty is large and declines as posterior means converge and posterior variance shrinks. Proposition 4 gives a separate asymptotic result. Under a unique efficient matching with a positive surplus margin, the posterior optimal matching coincides with the efficient matching after some finite time almost surely.

Signal precision affects posterior responsiveness. In the simulations, more precise signals produce somewhat more first period reallocation, while the date 20 overlap changes little. The convergence result in Proposition 4 also depends on the surplus margin separating the efficient matching from its alternatives.

Industry structure also affects worker payoff dispersion. In the simulations, stronger industry specific comparative advantage raises both within industry and cross industry dispersion, while the ratio between them changes little. The decomposition is specific to the worker optimal core used in the numerical exercise.

A further implication concerns the interaction of observable context and learned match quality. Early allocations are driven largely by the observable component u_{ij}, since posterior uncertainty about ε_{ij} is initially large. Over time, learning can overturn context-based rankings, especially when signals are precise (Proposition 2). Interpreted in labor market terms, this maps to a shift from selection on observables (credentials, firm reputation, industry affiliation) toward selection on realized performance.

Although the model is not a lifecycle framework, it also gives a simple account of wage dynamics. Entry wages reflect observable context, while subsequent wage adjustments reflect revisions in posterior expected productivity $v_{ij}(t) = u_{ij} + \mu_{ij,t}$. Workers whose realized signals exceed prior expectations experience faster wage growth and upward rematching pressure; workers whose signals disappoint experience slower growth and may separate if outside options improve under updated beliefs.

8.2. Empirical considerations

This subsection outlines empirical considerations and explains what would be required to implement the framework with data. The paper is primarily theoretical, so we do not develop identification or estimation results. With that scope in mind, the parameters $\{\beta, \delta, \xi, \lambda\}$ in the parametric example of Section 3.2 are of a form for which the cross-sectional identification arguments developed in the empirical matching literature on transferable utility models (Galichon and Salanié, 2022) are closely related. In practice, observing industry assignments $g(j)$ and exploiting within- and across-industry variation is the natural way to separate the industry complementarity component $\lambda\theta_{i,g(j)}$ from firm fixed effects. A rigorous identification and estimation treatment in this framework is left to future work.

Learning shows up only over time, so seeing it in data would require repeated observations of the same agents. In the model, the prior variance τ_0^2 sets the initial spread of unknown match quality, and σ^2 sets how fast posterior beliefs concentrate. Under the Normal-Normal specification, posterior variance evolves as $\tau_t^2 = \left(\tau_0^{-2} + t\sigma^{-2}\right)^{-1}$, so uncertainty about the idiosyncratic component shrinks at rate $O\left(\frac{1}{t}\right)$; in settings where wages and mobility are closely tied to posterior valuations, learning-driven variation should become less important over time, holding institutions fixed. Deviations from this pattern may indicate non-Gaussian learning, heterogeneity in signal precision, time varying evaluation regimes, or the presence of additional dynamics such as human capital accumulation or search frictions.

The model yields several qualitative predictions that can be brought to data. First, settings with more objective or standardized performance measurement (lower σ^2) should place greater weight on observed performance in posterior valuations. In our simulations, more precise signals are associated with somewhat higher first period reallocation. Second, in our simulations a stronger industry specific comparative advantage (higher λ) raises both within industry and cross industry payoff dispersion under the worker optimal core selection, without a clear shift in their ratio. Third, conditional on total expected surplus, separations should be more likely when posterior means deteriorate relative to context, since learning directly shifts the ranking of outside opportunities.

8.3. Connections to related empirical literatures

A large empirical literature studies employer learning about worker ability from performance observations, typically in environments where the incumbent employer observes richer information than outside firms (Altonji and Pierret, 2001; Lange, 2007). By contrast, the public signal benchmark here is most appropriate when performance information is verifiable and portable (e.g., third party evaluations, standardized tests, platform ratings), so that outside firms can condition offers on similar information. The framework also relates to empirical matching models that estimate multidimensional worker and firm heterogeneity from matched data (Bonhomme et al., 2019; Taber and Vejlin, 2020). Its contribution on this front is to separate observables from initially unobserved match quality that is progressively revealed, implying that the relationship between worker and firm components in wage decompositions can evolve over tenure. Finally, it complements work on dynamic sorting and search by isolating the learning channel in a frictionless assignment environment, which is most relevant in thick markets with low search costs (Eeckhout and Kircher, 2018).

The environments that motivate the public signal benchmark share a common feature. An external scoring technology produces verifiable information that the whole market can see, and this information is available before any particular match is formed. In academic job markets, publication records, citation counts, and letters of recommendation are visible to the entire market before departmental fit is learned on the job. In professional sports, combine metrics, draft eligibility evaluations, and minor league statistics are collected and disseminated independently of which team eventually signs the player. On online labor platforms, verifiable credentials and portable ratings accumulate across clients and are observable to future employers before any specific engagement. In medical residency and fellowship matching, board scores, clerkship grades, and standardized evaluations are visible market wide before on-the-job fit is revealed. In each case the decomposition $\pi = u + \varepsilon$ is natural: the scoring technology produces u for all potential pairs, while ε is revealed only gradually once a match is formed or performance is observed. The benchmark is less appropriate in settings where informative performance signals exist only within realized relationships; such environments are better captured by the matched pairs learning extension of Section 8.4.

Corblet et al. (2025) develop an empirical framework for repeated matching games in which the same population of agents rematches period after period, and the analyst estimates structural parameters from the resulting panel of match outcomes. The source of dynamics in that framework is the repeated game structure together with dynamics in fundamentals. Here, by contrast, dynamics come entirely from the evolution of posterior means under a fixed prior and a fixed signal technology; the solution concept at each date is static, and learning is exogenous to the matching sequence. The two approaches isolate different channels. Similarly, Doval and Schenone (2024) develop consistent conjectures as a dynamic stability notion, in which dynamics are driven by agents' beliefs about continuation play rather than by Bayesian updating about fixed unknowns.

8.4. Limitations and directions for future research

A central limitation of the benchmark is that signals are observed for all potential pairs independently of realized matches. As noted in Section 2.3, this is a demanding requirement, since observable evaluation systems typically record worker level ability or in match performance rather than pair level information about matches that never happened. In many labor market applications, informative performance signals are generated primarily through realized matches (or at least through interviews and trials), so the information set becomes endogenous to the matching sequence. Under matched pairs learning, pairs that are never formed accumulate no information, and a purely myopic surplus maximizing rule need not generate exploration of off path pairs. Establishing eventual selection of the efficient matching in that environment therefore typically requires additional ingredients (e.g., random meetings, forced experimentation, trial stages, or other mechanisms ensuring sufficiently rich sampling). We view that case as a distinct extension rather than a minor modification of the present benchmark. Beyond this decoupling, the public information benchmark also rules out several important forces. If incumbent firms observe worker performance more precisely than outside firms, separations can be adversely selected and poaching is disciplined by informational frictions (Greenwald, 2018; Waldman, 1984), while employer-specific learning can make mobility less informative about portable productivity (Gibbons and Katz, 1991). Allowing for private signals would therefore require a stability concept that accommodates heterogeneous beliefs, as in matching with incomplete information (Liu et al., 2014), and would introduce equilibrium selection issues that are absent under common posteriors.

The model also treats signal arrival as exogenous and abstracts from strategic behavior. Endogenous effort, endogenous evaluation intensity, and strategic timing of search would turn learning into an equilibrium object and connect the framework to dynamic contracting and mechanism design considerations. Incorporating lifecycle features such as retirement, discounting, or human capital accumulation would further enrich the dynamics but would require explicit state variables and an intertemporal payoff structure.

A further tension worth naming concerns the choice between structured and fully flexible correlation in the latent component. The industry specification of Section 3.2 makes the λ comparative static of Proposition 3 interpretable and empirically recognizable, but it does not exhaust the possible correlation structures across pairs. A fully flexible correlation specification would deliver greater generality at the cost of the closed-form posterior in Section 4, without sharpening Proposition 4. Section 4.6 offers a middle ground by relaxing pairwise independence through a component structure while preserving Gaussian conjugacy and Proposition 4(a).

Finally, tractability relies on the Normal-Normal structure. Non-Normal priors, heavy tailed shocks, or mis specified learning rules could alter the speed and pattern of convergence and would typically require numerical methods. Likewise, embedding the environment in general equilibrium, allowing participation, entry and exit, or education decisions to respond to expected match values, would generate feedback effects beyond the paper's partial equilibrium focus.

9. Conclusion

This paper develops a tractable framework for two sided matching markets with transferable utility in which agents learn about match quality through publicly observable signals. Match surplus combines an observable component, capturing worker and firm characteristics and industry structure, with an initially unknown idiosyncratic component that is revealed over time. The key simplification is that common posterior beliefs induce a valuation matrix at each date, and the market is analyzed as a static Shapley-Shubik assignment game on that matrix. Stability is defined as membership in the core of the sequence of static assignment games induced by current beliefs.

The analysis yields three main sets of results. First, under Normal-Normal learning, posterior means converge almost surely to true idiosyncratic match quality and posterior uncertainty shrinks at rate $O\left(\frac{1}{t}\right)$, providing a transparent characterization of how quickly information is aggregated. Second, signal precision governs responsiveness: higher precision increases the weight placed on realized performance in posterior valuations, strengthening the capacity of learning to overturn context-based rankings. Third, industry structure shapes cross sectional differentiation. Stronger industry complementarity amplifies differences in the observable component of surplus across industries, while within industry differences are driven mainly by firm level components. In the simulations, stronger industry complementarity raises both within industry and cross industry worker payoff dispersion.

The central long run implication is that information driven mismatch is transient under a mild margin condition. When the true surplus matrix admits a unique efficient matching separated by a positive margin, the posterior optimal matching coincides with the efficient matching after a finite random time almost surely, and once enough signals have arrived the probability of selecting a different matching at a given date admits an explicit exponential upper bound. In environments with verifiable and widely observed performance metrics, learning driven turnover is concentrated early and declines as uncertainty resolves.

Economically, the framework clarifies how the relative importance of credentials versus performance varies across settings. When performance signals are precise, realized outcomes rapidly discipline initial context-based beliefs; when signals are noisy, context remains influential for longer and early career sorting on observables persists.

The tractability of the approach makes it useful as a benchmark for empirical and theoretical work. In settings with verifiable performance measures, the model offers a way of thinking about turnover as a result of information arrival rather than of search frictions alone. At the same time, the benchmark points to natural extensions. Allowing for private signals would introduce adverse selection and belief heterogeneity and would require refined stability concepts; incorporating endogenous information acquisition or effort would make learning an equilibrium object; and embedding the framework in a lifecycle or general equilibrium environment would connect learning driven rematching to participation, entry and exit, and human capital investment.

The paper shows that combining Bayesian learning with the assignment game structure yields sharp, interpretable predictions about sorting, wages, and turnover in environments where performance information is widely observable. As measurement technologies and portable performance signals become more prevalent, understanding matching dynamics under learning becomes increasingly important for interpreting labor market outcomes and for evaluating policies that affect information transparency.

Funding

This research was financially supported by the French National Research Agency (ANR),"Investissements d'Avenir"(ANR-11-IDEX0003/LabEx Ecodec/ANR-11-LABX-0047), by the UK Research and Innovation (UKRI Frontier Research Grant EP/U537093/1), and by the European Union (SInfoNiA ERC-2023-ADG 101142530)" for Olivier Gossner's research.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

Data will be made available on request.

References

Altonji, J.G., Pierret, C.R., 2001. Employer learning and statistical discrimination. Q. J. Econ. 116 (1), 313–350.
Atay, A., Núñez, M., 2019. A note on the relationship between the core and stable sets in three-sided markets. Math. Soc. Sci. 98, 10–14.
Bikhchandani, S., 2017. Stability with one-sided incomplete information. J. Econ. Theory 168, 372–399.
Bonhomme, S., Lamadon, T., Manresa, E., 2019. A distributional framework for matched employer employee data. Econometrica 87 (3), 699–739.
Camiña, E., 2006. A generalized assignment game. Math. Soc. Sci. 52 (2), 152–161.
Chen, Y.C., Hu, G., 2020. Learning by matching. Theor. Econ. 15 (1), 29–56.
Choo, E., Siow, A., 2006. Who marries whom and why. J. Polit. Econ. 114 (1), 175–201.
Corblet, P., Fox, J.T., & Galichon, A. (2025). *Repeated matching games an empirical framework*.
Doval, L., Schenone, P., 2024. Consistent conjectures in dynamic matching markets. Math. Soc. Sci. 132, 114–127.
Eeckhout, J., Kircher, P., 2018. Assortative matching with large firms. Econometrica 86 (1), 85–132.
Forges, F., 2004. The ex ante incentive compatible core of the assignment game. Math. Soc. Sci. 47 (2), 135–151.
Gale, D., Shapley, L.S., 1962. College admissions and the stability of marriage. Am. Math. Mon. 69 (1), 9–15.
Galichon, A., Salanié, B., 2022. Cupid's invisible hand: social surplus and identification in matching models. Rev. Econ. Stud. 89 (5), 2600–2629.
Gibbons, R., Katz, L.F., 1991. Layoffs and lemons. J. Labor. Econ. 9 (4), 351–380.
Greenwald, B.C., 2018. Adverse Selection in the Labor Market. Routledge.
Jovanovic, B., 1979. Job matching and the theory of turnover. J. Polit. Econ. 87 (5, Part 1), 972–990.
Lange, F., 2007. The speed of employer learning. J. Labor. Econ. 25 (1), 1–35.
Liu, Q., Mailath, G.J., Postlewaite, A., Samuelson, L., 2014. Stable matching with incomplete information. Econometrica 82 (2), 541–587.
Qi, D., 2025. The rates of learning with public and private signals. Math. Soc. Sci. 136, 102433.

Roth, A.E., Sotomayor, M., 1992. Two-sided matching. Handb. Game Theory Econ. Appl. 1, 485–541.

Shapley, L.S., Shubik, M., 1971. The assignment game I: the core. Int. J. Game Theory 1 (1), 111–130.

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

Waldman, M., 1984. Job assignments, signalling, and efficiency. Rand J. Econ. 15 (2), 255.