I Z A Institute of Labor Economics
Initiated by Deutsche Post Foundation

DISCUSSION PAPER SERIES

IZA DP No. 16986

Stable Matching on the Job? Theory and Evidence on Internal Talent Markets

Bo Cowgill
Jonathan M.V. Davis
B. Pablo Montagnes
Patryk Perkowski

MAY 2024

I Z A Institute of Labor Economics

Initiated by Deutsche Post Foundation

DISCUSSION PAPER SERIES

IZA DP No. 16986

Stable Matching on the Job? Theory and Evidence on Internal Talent Markets

Bo Cowgill
Columbia Business School and IZA

Jonathan M.V. Davis
University of Oregon

B. Pablo Montagnes
Emory University

Patryk Perkowski
Yeshiva University

MAY 2024

ISSN: 2365-9793

IZA – Institute of Labor Economics

Schaumburg-Lippe-Straße 5–9
53113 Bonn, Germany

Phone: +49-228-3894-0
Email: publications@iza.org

www.iza.org

ABSTRACT

Stable Matching on the Job? Theory and Evidence on Internal Talent Markets*

A principal often needs to match agents to perform coordinated tasks, but agents can quit or slack off if they dislike their match. We study two prevalent approaches for matching within organizations: Centralized assignment by firm leaders and self-organization through market-like mechanisms. We provide a formal model of the strengths and weaknesses of both methods under different settings, incentives, and production technologies. The model highlights tradeoffs between match-specific productivity and job satisfaction. We then measure these tradeoffs with data from a large organization's internal talent market. Firm-dictated matches are 33% more valuable than randomly assigned matches within job categories (using the firm's preferred metric of quality). By contrast, preference-based matches (using deferred acceptance) are only 5% better than random but are ranked (on average) about 38 percentiles higher by the workforce. The self-organized match is positively assortative and helps workers grow new skills; the firm's preferred match is negatively assortative and harvests existing expertise.

JEL Classification: M5, D47, J4

Keywords: internal labor markets, assortative matching, assignment mechanisms, team formation, matching

Corresponding author:
Bo Cowgill
Columbia Business School
665 W 130th St
New York, NY 10027
USA

E-mail: bo.cowgill.work@gmail.com

* The authors thank participants at NBER Economics of Organizations (2022), NBER/CEME Decentralization (2022), SIOE, Wharton People and Organizations, the Marketplace Innovations Workshop, AOM, the Workshop for Information Systems and Economics (WISE), SMS, and the 2021 Causal Data Science Meeting, and seminar participants at Carnegie Mellon, Columbia, Cornell, MIT, Rochester, Stanford (GSB), University of Minnesota, and University of Hong Kong. We also thank Amanda Dahlstrand, Bob Gibbons, John Hatfield, Navin Kartik, Yash Kanoria, Fuhito Kojima, Scott Kominers, Jin Li, John Morgan, Parag Pathak, Alex Teytelboym and Brian Wu for helpful comments. Cowgill thanks the Kauffman Foundation. Earlier versions of this paper were called, "Matching for Strategic Organizations," "Preferences and Productivity in Job Matching," "Designing Organizational Versus Public Markets," and "Matchmaking Principals: Theory and Evidence from Internal Labor Markets."

1 Introduction

How do organizational leaders match personnel with managers, co-workers, and tasks? Unhappy workers could quit or become demotivated. However, pleasing workers is not most leaders' main objective. Left to their own devices, workers and managers may pursue their own goals without regard for the overall organization's objectives and strategy.

In this paper, we study how different allocation methods shape organizational matches, teams, and their performance on company objectives. While there is a large literature on misallocation between firms (e.g. Hsieh and Klenow, 2009), less is known about misallocation between workers and jobs within firms (Benson et al., 2019, 2021; Haegele, 2021; Friebel and Raith, 2022), and how various matching methods can impact organizational outcomes.

We study two widely-used approaches to allocating workers. On one hand, leaders may match workers to jobs by fiat ("firm-dictated") based on their own information and objectives. This avoids bad matches from the principal's perspective, e.g., matches where workers lack key skills or qualifications.

On the other hand, firms can open internal markets where workers and managers self-organize assignments. By catering to workforce preferences, "internal talent markets" aim to increase engagement, job satisfaction, and retention. Some of these systems explicitly borrow mechanisms from the market design literature.[1] Internal talent markets are part of a multi-decade trend in which responsibility for talent management has shifted from central human resources (HR) to the frontline workforce.[2] Organizations such as

[1] We review a list of organizations using market design tools for internal talent markets in Section 2. The list includes Google (Cowgill and Koning, 2018), the US Army (Davis et al., 2023), Teach For America (Davis, 2022), and the International Monetary Fund (Barron and Vardy, 2005). Vardy confirmed in private correspondence the IMF proposal was implemented.

[2] This trend is described by Cappelli (2013) and documented by Whittaker and Marchington (2003), Perry and Kulik (2008), Friebel and Raith (2022), and others.

Google, Wal-Mart, Accenture, and the United States Army have adopted market-based systems for internal assignments. An ecosystem of venture-backed startups provides software for internal markets,[3] with clients around the Fortune 500 and a market capitalization of over $1B.[4]

We examine how these two allocative methods drive the efficiency of matching inside organizations. The key trade-off we study is the balance between match value and job satisfaction (particularly retention). Centralized alignment ensures that workers are matched to jobs where they will be productive, but not necessarily where they will be happy. By contrast, delegating to markets allows workers to self-select into the jobs they like (and thus improves retention and motivation), but possibly at the expense of match-specific productivity. Although we focus on companies' internal organization, there are parallel issues in other domains.

We develop a formal model of the strengths and weaknesses of either allocation method. Dictating requires leaders to be accurately-informed about match-specific productivities; otherwise, internal markets are preferable. However, even if leaders are well informed, the merits of dictating still hinge on the workforce's alignment with leadership's goals and key features of the business environment.

We show that if a workforce's preferences are *not* aligned (no correlation with leadership's goals), the principal will prefer dictating i) if match-specific productivity is high, or ii) if match-specific preferences are weak. By contrast, even if a workforce's preferences *are* aligned with the principal's, markets are not necessarily better. Leaders can sometimes achieve the same assignments by fiat, without compromising job satisfaction. Coordination problems (multiple equilibria) also appear between aligned workers. Even

[3]These startups include Gloat (http://gloat.com), Fuel 50 (http://fuel50.com), and Hitch Works (https://hitch.works/). Avela (https://avela.org/) is another social impact software and consulting company with academic ties which offers marketplace software to organizations.

[4]Gloat's clients include Wal-Mart, PepsiCo, Vanguard, Unilver, and Nestlé. For Gloat's valuation, see https://techcrunch.com/2022/06/28/gloat-nabs-90m-to-build-ai-powered-internal-jobs-marketplaces/.

with highly-aligned workers, firms may prefer to dictate matching because of the coordination benefits.

Our empirical section shows that the tradeoffs between these two approaches are large in practice. We compare market-driven assignments to those chosen by the firm's executives using detailed data from a large organization. Our research design uses a match-specific scoring algorithm developed by the firm to assess match quality (using characteristics of workers and jobs).[5] We combine this with preference data submitted by the workforce to the deferred acceptance algorithm ("DA") that was used to match workers to teams at the firm.[6] We assess the executives' view of DA-generated matches using the firm's own scoring rule. We then contrast the DA matches to alternative assignments that maximize the firm's objectives (using the match quality scores). We assess these "firm-preferred" matches through the workforce's lens using the preference data submitted to DA.

We find large differences between mechanisms, in both the business value of matches and workforce satisfaction. Using the firm's match quality score, our results suggest that the firm-dictated match is 33% more valuable than randomly assigned matches within job categories. This is a multiple-standard deviation improvement: To achieve the same improvement by training or replacement, the bottom 80% of workers must perform as well as the 80th percentile employee. By contrast, the quality of job assignments created through delegation are only 5% better than random matching. Thus, there are large match quality benefits to dictating assignments in our setting.

While dictating assignments has large match quality benefits, delegating assignments performs much better on measures of worker satisfaction. Through deferred acceptance,

[5]We describe the development and training of the match quality model in Section 5.

[6]For readers unfamiliar with the deferred acceptance algorithm (Gale and Shapley, 1962; Roth and Sotomayor, 1990), Appendix **??** provides a description. A key feature of DA is that it is strategyproof for the proposing side; in our setting, the dominant strategy for the workers was to submit their honest rankings of managers.

workers and managers rank their assignments better by an average of about 38 percentiles in rank. They are 75 percentage points less likely to be assigned to a role they ranked as "tied for last place" (but as better than quitting). By contrast, the workforce ranks the firm's matches as approximately as good (or bad) as randomly-chosen assignments. Thus, the assignment mechanisms also feature large trade-offs on worker satisfaction.

Why do these differences arise? We show that match-specific productivity (as defined by executives) is not a significant driver of workforce preferences. Instead, workforce preferences differ along two key dimensions. First, the firm-dictated match is more *negatively* assortative: In this allocation, the best managers are not necessarily matched with the best workers.[7] By contrast, the workforce's preferences generate more *positively* assortative matching: high-quality workers and high-quality managers are paired together. In addition, workers prioritize roles that allow growth in skills they desire to improve. By contrast, executives prefer that workers perform well immediately in their match (without on-the-job training). A prior literature explains why employers avoid training workers, particularly for non-firm specific skills (poaching externalities, Becker, 1964). In our setting, about 90% of skills that workers prioritized in the market were *non-* firm-specific.

Had participants sought to maximize their own match quality in their next job — ignoring the externalities this would impose — then deferred acceptance would have resulted in a 24% increase above random assignments. To achieve the full 33% increase, some participants would need to internalize the externalities of their choices.[8] We find no evidence of such internalization. We also find little evidence that the workforce actually shared the CEO's goals, but simply failed to coordinate on the best (among the multiple) equilibria.

Despite these tradeoffs, we cannot characterize either approach (firm-dictated or market driven) as a mistake in our setting. Although the workforce does not appear to be aligned,

[7]The firm's preference appears to be driven by opportunity cost considerations: Assigning a star worker to a top manager "wastes" the worker contributions on projects that easily succeed anyway. For other papers about negative assortative matching, see Becker (1973); Lazear et al. (2015); Adhvaryu et al. (2020).

[8]For example, by accepting lower quality matches for the benefit of the firm as a whole.

the internal market can still be justified on retention grounds. Both approaches have significant costs and benefits.

The two matching approaches in this paper can be viewed as extremes. Firm-dictated uses the information available to the CEO and not the workforce (information about match-specific contributions to firm goals). Market-like delegation uses information available to the workforce and not the CEO (agents' private match-specific preferences). Our paper studies the gravitational forces pulling organizations closer to one extreme (or the other) and how large these forces are in a real-life setting. Although we do not study approaches between the extremes, they are a promising avenue for future research. Our model suggests tradeoffs for matching mechanisms between the poles. As the CEO inches towards delegating, the organization may enjoy greater ability to incorporate distributed information. However, the CEO loses control and workers may use their discretion for private gains (at the CEO's expense).

The remainder of this paper is organized as follows. Section 2 discusses related literature. Section 3 presents a model of organizational matching with participation constraints. Sections 4 through 6 review our empirical setting, strategy and main results. Section 7 discusses reasons for the misalignment and Section 8 contains an empirical extension about the level of CEO knowledge. Section 9 concludes.

2 Related Literature

Coordination within organizations. We model endogenous team formation (matching) as a type of decentralized management structure (Arya et al., 2002; Ortega, 2001; Hamilton et al., 2003; Christensen and Knudsen, 2010; Allocca, 2023).

This problem has similarities with other principal/agent problems; for example, the choice to delegate in the presence of biased, informed agents (Dessein, 2002; Alonso and

Matouschek, 2008). The multi-agent, two-sided nature of the problem requires coordination, even when all agents are "aligned" with their principal.

We connect this literature to the research about "happy workers being more productive" (Oswald et al., 2015; Bellet et al., 2023). A closely related literature examines non-monetary incentives, for example, workers who find intrinsic motivation or meaning in their assignments (Cassar and Meier, 2018). There is a growing interest in understanding matching managers and workers to positions. Firms may think differently about this than their workers and managers (Haegele, 2021; Cestone et al., 2022; Friebel and Raith, 2022). Keeping workers happy and motivated is not necessarily free to the firm. In Bandiera et al. (2010) and Park (2019), workers' productivity is correlated with the social contact between co-workers. In Xu et al. (2021), civil servants prefer to be matched with jobs in their home states, but are judged as less effective when they are. Our model places these considerations at the center of firms' resource allocation.

Our paper also links research about information technology (IT) and organizations (Brynjolfsson and Hitt, 2000; Bresnahan et al., 2002; Bloom and Van Reenen, 2011). Several results in our theory can be seen as modeling IT quality. Routine-biased technical change (RBTC) could increase employer demand for "generalists," and influence design of internal talent markets. Our results contribute to the understanding of IT, RBTC, and organizational structure (Lindbeck and Snower, 2000; Caroli and Van Reenen, 2001; Bloom and Van Reenen, 2011).[9]

Stable Matching. In the stable matching literature (Roth and Sotomayor, 1990), agents can often match outside a centralized mechanism if they prefer. As a result, market de-

[9]Self-organization has also attracted significant attention from management scholars (Raveendran et al., 2021; Lee and Edmondson, 2017) based partly on case studies of Valve, Zappos, and Morning Star (Martela, 2019). Our paper is conceptually related to these management practices, although we focus on a particular stylized implementation.

signers often restrict attention to mechanisms that generate pairwise stable matches.[10] As we discuss in Section 3, organizational leaders can often block matches. While organizational leaders have incentives to consider workforce preferences (because workers can choose to quit), the ability to block matches allows CEOs to address them without pairwise stability constraints.

Market allocation and policy goals. A smaller literature studies matching mechanisms to achieve policy goals (Abdulkadiroglu et al., 2021; Combe et al., 2022; Thakur, 2021). Several empirical papers in this literature examine counterfactuals in which allocations are dictated by a planner, rather than through a market-like mechanism (Agarwal et al., 2020; Ba et al., 2021; Dahlstrand, 2021; Bates et al., 2022). We model this difference, and contribute an empirical personnel application.

Constraints (such as quotas) are a major design tool for pursuing policy goals. Kojima et al. (2020) show that constraints are particularly useful when a firm's production technology is *group separable*. Group separability requires divisions within an organization to not compete against each other for talent, and for division leaders to be aligned with the CEO. Neither condition is met in many real-world organizations. In our model of firm-dictated assignment, a key benefit is the ability to coordinate competing divisions.

Internal talent markets. Although our model is more general, internal talent markets are our leading example (Baker and Holmstrom, 1995). Markets explicitly using market design tools have appeared at Google (Cowgill and Koning, 2018), Teach For America (Davis, 2022), the US Army (Sonmez and Switzer, 2013; Davis et al., 2023), the International Monetary Fund (Barron and Vardy, 2005), the UN World Food Programme (Delgado, 2014), and others.[11] In July 2021, a report by Georgetown and Harvard Universi-

[10] In this context, pairwise stable means no unmatched pair of agents would prefer to be matched together over their assigned match.

[11] Other examples include Cowgill et al. (2020), which studies a large Asian bank with "millions of customers, billions of dollars in assets and in revenues, and thousands of employees" that uses deferred accep-

ties indicated that the US State Department is collaborating with the National Resident Matching Program (NRMP) to develop a system for assigning diplomatic officers to overseas posts.[12]

Firms such as McKinsey and Deloitte have adopted formal internal talent marketplaces (ITMs) and advocated them for clients (sometimes in collaboration with the aforementioned software vendors.[13] Although the microstructures of these marketplaces are not disclosed in detail, their descriptions show that respecting participants' preferences (the key property in our theoretical results) is a guiding principle.

3 Theoretical Framework

Below we model a firm's choice either to dictate matches or to delegate matching to an internal market. Proofs for the results are in Appendix **??**.

Preliminaries. An organization consists of a principal (e.g., a CEO) and two types of agents: workers and divisions (led by middle managers). We refer to these groups collectively as "agents," or "the workforce." In general, there are I workers $i \in \{1, \ldots, I\}$ and J divisions $j \in \{1, \ldots, J\}$. We focus attention on balanced markets where $I = J = N$. All agents can be matched with a single member of the other side of the market. There are N ways to match each agent, and $N!$ ways of matching all agents in an $N \times N$ matching. Each $N \times N$ matching can be indexed by $a \in A$, with $|A| = N!$. $\alpha_{ij}(a)$ is a match indicator

tance to match and re-allocate workers, managers and projects. The United States Military Academy uses a cumulative offer mechanism to assign cadets to branches (Sonmez and Switzer, 2013). The International Monetary Fund used the deferred acceptance algorithm to assign new economists to research teams (Barron and Vardy, 2005).

[12] See https://isd.georgetown.edu/2021/07/02/new-research-highlights-retention-crisis-at-the-state-department.

[13] For McKinsey, see https://www.mckinsey.com/business-functions/organization/our-insights/making-a-market-in-talent. For Deloitte, see https://www2.deloitte.com/content/dam/insights/us/articles/4582_are-you-overlooking-your-greatest-source-of-talent/DI_are-you-overlooking-your-greatest-source-of-talent.pdf.

that is equal to 1 if worker i and division j are matched in assignment a, and 0 otherwise.

CEO Preferences. The CEO has a value v_{ij} for a pairwise match between i and j. For a private company, v_{ij} could represent the match's contribution to total firm profits, and for non-profits or governments v_{ij} could represent social objectives. One could think of the v_{ij}s as *valuations* akin to a bidder's valuation for an object in an auction.

Definition 1 (CEO Value for Allocations)**.** *The CEO's value for allocation a is $V(a)$, or the sum of all productivities v_{ij} that are assigned in the matching a:*

$$V(a) = \sum_{i=1}^{I} \sum_{j=1}^{J} \alpha_{ij}(a) v_{ij}. \tag{1}$$

We next create a measure of how sensitive a firm's production function is to matching.

Definition 2 (Distribution of V)**.** *The function $\mathcal{V} : V(a) \mapsto \mathbb{N}$ is the frequency distribution of $V(a)$.*

$\mathcal{V}(x)$ expresses how many allocations a have the same $V(a) = x$. A firm whose $\mathcal{V}$ places all mass on one point will always have the same output, as long as nobody quits, irrespective of how workers are assigned. $\mathcal{V}$s with wider support represent firms whose output depends heavily on matching. The properties of $\mathcal{V}$ will later affect the CEO's choice of mechanism.

Workforce Preferences. Without any other constraints, the CEO can maximize her utility by selecting the allocation that maximizes $V(a)$ (the sum of all assigned v_{ij}s). However, workers and divisions have preferences over their match partners. There are many ways a CEO could be penalized for ignoring these preferences. Below we develop a model of quitting. However, unhappy agents could also become demotivated, exert lower effort, demand higher wages, or create other negative payoffs to the CEO. In some firms, information flows and knowledge-sharing are crucial in supporting effective coordination (Impink et al., 2021), and unhappy workers might share less information. For these reasons, the CEO may consider incorporating the workforce's preferences into matching.

To formalize the workforce's preferences, we say that worker i's utility from matching with division j is μ_j^i (and division j's utility from this match is u_i^j). To help conceptualize the variables, Figure 1 presents an example of two workers and two divisions, showing the match-specific valuations for all players (workers, division managers and the CEO).

Figure 1: **Match-Specific Valuations for All Players (2 $\times$ 2 Example)**

	Division 1	**Division 2**
Worker 1	**W1** : μ_1^1 **D1** : u_1^1 **CEO** : v_1^1	**W1** : μ_2^1 **D2** : u_1^2 **CEO** : v_2^1
Worker 2	**W2** : μ_1^2 **D1** : u_2^1 **CEO** : v_1^2	**W2** : μ_2^2 **D2** : u_2^2 **CEO** : v_2^2

Worker Valuations: μ_j^i, **Division Valuations:** u_i^j, **CEO Valuations:** v_j^i

Notes: This figure displays the match specific valuations for all players in a simple 2 $\times$ 2 version of our model. This includes two workers, two divisions and one CEO.

We abstract away from the agent's reasons behind the preferences. Some workers may like the manager interpersonally, or they may like the manager's project. Workers may also be drawn to a manager because of the manager's existing team members. Similarly, managers' preferences over workers may be driven by the worker's skills, personality or fit with her existing team. Managers could anticipate peer effects (Mas and Moretti, 2009), or information sharing between a potential worker and the incumbent team (Alonso and Matouschek, 2008; Allocca, 2023). We now make two substantive assumptions.

Assumption 1 (Quits Reduce Output)**.** *A match is acceptable to a worker or division if it yields greater utility than their outside option* $\underline{u}$*. If either the worker or the division finds the match* *unacceptable, the agent quits and their abandoned match yields zero output for the CEO.*

Because of Assumption 1, the CEO might not consume v_{ij} even if she assigns $i \leftrightarrow j$.

Assumption 2 (No Match-Specific Transfers)**.** *The CEO cannot make match-specific transfers.*

In many practical settings, wages are inflexible. This includes the government and/or unionized jobs.[14] Even in private sector, non-unionized settings, laws or fairness expectations may limit pay flexibility.[15] Our model therefore rules out match-specific pay; however, future work could incorporate it for settings where it is allowed. Using Assumptions 1 and 2, we write the CEO's payoff for selecting a given allocation a.

Definition 3 (CEO Payoffs)**.**

$$\pi(a) = \sum_{i=1}^{I} \sum_{j=1}^{J} \alpha_{ij}(a) \cdot \underbrace{v_{ij}}_{\substack{\text{Value} \\ \text{to CEO}}} \cdot \underbrace{\mathbb{1}(\mu_j^i \geq \underline{u})}_{\substack{\text{Acceptable} \\ \text{to Worker?}}} \cdot \underbrace{\mathbb{1}(u_i^j \geq \underline{u})}_{\substack{\text{Acceptable} \\ \text{to Manager?}}} \tag{2}$$

Table 2 shows CEO payoffs in our 2×2 example from Figure 1.

Table 1: **CEO Payoffs (2×2 Example)**

π(On Diagonal)	$v_1^1 \cdot \mathbb{1}(u_1^1 \geq \underline{u}) \cdot \mathbb{1}(\mu_1^1 \geq \underline{u}) + v_2^2 \cdot \mathbb{1}(u_2^2 \geq \underline{u}) \cdot \mathbb{1}(\mu_2^2 \geq \underline{u})$
π(Off Diagonal)	$v_2^1 \cdot \mathbb{1}(u_2^1 \geq \underline{u}) \cdot \mathbb{1}(\mu_1^2 \geq \underline{u}) + v_1^2 \cdot \mathbb{1}(u_1^2 \geq \underline{u}) \cdot \mathbb{1}(\mu_2^1 \geq \underline{u})$

Notes: This table shows the payoffs to the CEO in the 2×2 example in Figure 1.

Finally, the CEO cannot necessarily pick the allocation a with the highest payoff because she may not know the workers' and managers' preferences.

[14]In unionized workforces and government jobs, pay scales do not permit assignment-specific pay. See, for example, https://www.dfas.mil/militarymembers/payentitlements/Pay-Tables/ for the US Army and State Department.

[15]The case study of Google's internal market states that pay raises "were legally required to go through Google's centralized HR system," rather than be settled by the internal market. One law, California's Fair Pay Act requires "Requiring equal pay for employees who perform 'substantially similar work,' when viewed as a composite of skill, effort, and responsibility'.' https://www.dir.ca.gov/dlse/california_equal_pay_act.htm.

Assumption 3 (Private Information). *Agents privately know their own preferences (μ_j^i or u_i^j) and outside options ($\underline{u}$). The CEO does not know the agents' preferences, but knows the distributions from which they were drawn. The CEO also knows the productivity of all pairs of workers and managers (v_{ij}) and the agents' outside option $\underline{u}$.*

The workforce's private information captures the idea that useful knowledge is often distributed throughout an organization. A key benefit of delegating comes from tapping this distributed knowledge. Using the information in Assumption 3, the CEO can calculate the probability that an $i \leftrightarrow j$ match is acceptable to both i and j, which we call $P(i \leftrightarrow j \text{ acceptable})$. We will later place a functional form on this term based on how workforce and CEO preferences align. The expected retention rate across all matches is shown below.

Definition 4 (Expected Retention).

$$\mathbb{E}[R(a)] = \frac{1}{N}\sum_{i=1}^{I}\sum_{j=1}^{J}\alpha_{ij}(a)P(i \leftrightarrow j\ acceptable) \tag{3}$$

Accounting for retention expectations, the CEO's expected payoffs are:

Definition 5 (Expected CEO Payoffs).

$$\mathbb{E}[\pi(a)] = \sum_{i=1}^{I}\sum_{j=1}^{J}\alpha_{ij}(a)v_{ij}P(i \leftrightarrow j\ acceptable) \tag{4}$$

Table 2 shows the CEO's expected payoffs in our 2×2 example from Figure 1.

Table 2: **Expected CEO Payoffs (2×2 Example)**

$\mathbb{E}[\pi(\text{On Diagonal})]$	$v_1^1 \cdot P(1 \leftrightarrow 1 \text{ acceptable}) + v_2^2 \cdot P(2 \leftrightarrow 2 \text{ acceptable})$
$\mathbb{E}[\pi(\text{Off Diagonal})]$	$v_1^2 \cdot P(1 \leftrightarrow 2 \text{ acceptable}) + v_2^1 \cdot P(2 \leftrightarrow 1 \text{ acceptable})$

Notes: This table shows the payoffs to the CEO in the 2×2 example in Figure 1.

3.1 The CEO's Problem

The risk-neutral CEO must choose a mechanism $\mathcal{M}$ to select an allocation $a_{\mathcal{M}}$ that maximizes her expected payoffs $\mathbb{E}[\pi(a_{\mathcal{M}})]$. That is:

$$\begin{aligned} \max_{\mathcal{M}} \quad & \underbrace{\sum_{i=1}^{I}\sum_{j=1}^{J}\alpha_{ij}(a_{\mathcal{M}})v_{ij}P(i \leftrightarrow j \text{ acceptable})}_{\mathbb{E}[\pi(a_{\mathcal{M}})],\ \text{Definition 5}}, \\ \text{s.t.:} \quad & \sum_{i=1}^{I}\alpha_{ij}(a_{\mathcal{M}}) = 1, \qquad \sum_{j=1}^{J}\alpha_{ij}(a_{\mathcal{M}}) = 1. \end{aligned} \tag{5}$$

The CEO's problem differs from the classic "assignment problem" (Kuhn, 1955) because of asymmetric information. However, the CEO could select the match with the highest expected payoff accounting for both productivity and retention differences. Because of the lack of input from workers and divisions, we call this mechanism "firm-dictated assignment" ("FD").

The CEO could also gather workers' preferences. In other settings, researchers have developed strategy-proof mechanisms for eliciting match preferences (Roth and Sotomayor, 1990). However, the CEO's problem is distinct from the market designer's. The CEO has preferences over the assignment. In other stable matching settings, the match institution is broadly indifferent towards outcomes (subject to stability constraints).[16] Second, although CEOs cannot stop workers from exiting, they do have the power to block unwanted matches (even if both worker and division care to proceed). Although the CEO cares about retention, the match does not need to be pairwise stable.

[16]For example, the National Medical Residency Match does not attempt to match agents according to its own views about which doctors should work where, but rather in a way that pleases doctors and hospitals (subject to stability constraints).

3.2 Structure of Matching and Preferences

We now study the generic problem above by introducing additional structure. We begin by making workforce preferences independent of CEO goals, and later make them interdependent.

Assumption 4 (Private Values)**.** *Workers' utilities* μ_j^i *and division utilities* u_i^j *for each match partner are independently drawn from a common, absolutely continuous distribution* G*, and a match is acceptable to each side with positive probability (*$G(\underline{u}) < 1$*).*

Misalignment between principals and agents with private values is common in economics of organizations. Assumption 4 eliminates interdependence between the CEO's objectives and the workforce's preferences. I.I.D. random preferences are also common in models of matching markets (Knuth, 1976; Pittel, 1992; Ashlagi et al., 2017).

Assumptions 1 and 4 makes each agent's preferences uncorrelated with their outside offer. In Appendix **??**, we show that our main results are robust to alternative assumptions in which agents have both I.I.D. random preferences and I.I.D. outside options. In Appendix **??**, we also allow outside offers to be correlated with productivity.[17]

Structure of Assignment Mechanisms. To implement "firm-dictated assignment," we show in Appendix **??** that FD is an application of the Kuhn-Munkres method (the "Hungarian algorithm") that incorporates quit probabilities (alongside v_{ij}). To implement delegation, several mechanisms exist for matching with preferences. We study one in particular:

Assumption 5 (Deferred Acceptance)**.** *The CEO implements delegated matching using the worker-proposing deferred acceptance algorithm (DA, Gale and Shapley, 1962; Roth, 2008).*

[17] Assumptions 1 and 4 also state that the two sides' preference are symmetric: both are drawn I.I.D. from the same distribution G, and both sides have a common outside option $\underline{u}$. This helps minimize notation and describe the important features of agent preferences (on either side) succinctly. The symmetry does not drive any of our results. In reality, preferences (and outside options) on either side are likely to be different.

DA is widely used in practice (Roth, 2008), including in the empirical setting of this paper (Section 4). DA provides a tractable model of delegated matches that is strategyproof for the proposing side (and both sides in some models of large markets).[18] Our results will be framed in terms of workers-proposing DA, but we highlight where proposing matters. Many of our results are likely more general than DA and may apply to a wider set of preference-driven mechanisms. We highlight where this is likely. For readers unfamiliar with DA, Appendix **??** provides an overview.

3.3 Delegation vs Firm-Dictated Assignment

Using the structure above, we can now compare the CEO's payoffs.

Lemma 1 (CEO Payoffs from Dictating)**.** *The CEO's expected payoff from dictating,* $\mathbb{E}[\pi(a_{FD})]$*, equals* $\bar{R}V(a_{FD})$*, where* $\bar{R}$ *is the average retention rate of all assignments (*$\bar{R} = (1 - G(\underline{u}))^2$*).*

In Lemma 1, a_{FD} is the allocation that maximizes $V(a)$. For brevity we use V_{FD} for $V(a_{FD})$ moving forward.

Lemma 2 (CEO Payoffs from Deferred Acceptance)**.** *The CEO's expected payoff from DA,* $\mathbb{E}[\pi(a_{DA})]$*, equals* $R_{DA}\bar{V}$*, where:*

1) *$\bar{V}$ is the mean $V(a)$ of all possible allocations* $\left(\bar{V} = \frac{1}{N!}\sum_{a\in A} V(a)\right)$*, and*
2) *R_{DA} is the expected retention rate of DA.*

Lemma 2 uses the fact that DA selects each matching a with equal probability (under the I.I.D. assumption). As a result, the CEO's expected $V(\cdot)$ of the DA match is $\bar{V}$, the average output over all $N!$ possible matchings. The full payoff of DA also depends on retention.

[18] See e.g., Immorlica and Mahdian (2005); Kojima and Pathak (2009).

Lemma 3. *Given Assumption 4, $R_{DA} \geq \bar{R}$. The expected retention rate of DA is higher than the average retention rate.*

Lemma 3 likely extends to other matching algorithms (besides DA) based on workforce preferences. To improve retention above $\bar{R}$, a matching approach simply needs to eliminate at least one undesirable match for workers or managers.

Delegation vs. Firm-Dictated Assignment. Together, these results imply FD is expected to yield greater output than DA when $\bar{R}V_{FD} \geq R_{DA}\bar{V}$. Rearranging, we can see that CEO's choice of FD or DA depends on whether improvements in productivity or retention are more important:

$$\frac{V_{FD}}{\bar{V}} \geq \frac{R_{DA}}{\bar{R}}. \tag{6}$$

The left-hand side is a property of $\mathcal{V}$, the firm's output technology; it essentially measures how sensitive a firm's output is to changes in matching. The right-hand side is a function of G, the distribution from which workers' and managers' preferences are drawn. Using these, we can derive our first set of results.

3.4 Match-Specific Output: The Benefit of Firm-Dictated Assignment

From the CEO's perspective, the benefit of FD is that output — conditional on retention — is higher. We begin by exploring the implications of this insight.

Definition 6 (Specialization)**.** *A workforce is* unspecialized *if the outputs (v_{ij}'s) for each worker i are equal for all possible assignments j. The workforce is* specialized *if workers' outputs vary across assignments.*

Our notion of specialization refers to match-specific output *within a firm's boundaries.*

Bloesch et al. (2022) calls specialization "position specific skills."[19] In an unspecialized workforce, all the mass in $\mathcal{V}$ is concentrated at a single point, and the LHS of Equation 6 is equal to 1. Specialization (or lack thereof) is a property of the production function and is not affected by the choice of assignment mechanism.

Proposition 1. *The performance of delegation will equal or exceed that of firm-dictated assignment in firms where the workforce is completely unspecialized.*

The intuition behind Proposition 1 is that rearranging unspecialized workers generates no productivity benefits for the CEO. If workers are equally productive in all divisions, then the output of the most productive match (V_{FD}) is equal to that of the average match ($\bar{V}$). By contrast, rearranging an unspecialized workforce could generate *retention* benefits. As long as baseline quit rates $G(\underline{u})$ are above zero, the retention rate of DA will exceed that of FD. For these reasons, an unspecialized workforce makes delegation more attractive for the CEO.

Corollary 1. *For firm-dictated assignment to outperform delegation, it is necessary (but not sufficient) for the workforce to be specialized.*

Eq. 6 shows why specialization alone is not sufficient. Specialization must also generate enough extra output to justify the retention penalty of firm-dictated assignment.

Extension 1 (Noisy CEO Beliefs) Until now, we have assumed that CEOs have perfect knowledge of workers' match-specific productivities (v_{ij}). In Appendix **??**, we consider an extension in which the CEO observes these *noisily*. Even if workers are specialized, the CEO may be unable to observe the specializations clearly.

Proposition **??** shows that noisy measurement diminishes the benefits of firm-dictated

[19] In theory, a "specialist" in our setting can perform some tasks better than others within a firm's boundaries – but could possibly do many others across the broader economy. Similarly, a "generalist" in our setting — such as a software engineer who could work on many internal projects in a tech firm — could be a "specialist" in the broader economy (limited only to software projects).

assignment. If the CEO does not know which matches are productive, then the firm is better off enjoying the retention benefits of delegating. Proposition **??** suggests that better monitoring and/or analysis technology can tilt organizations towards central planning.

3.5 Retention: The Benefit of Delegation

Here, we explore the potential retention benefits of DA.

Proposition 2. *The retention benefits of DA relative to FD are higher as the unconditional quit probability $G(\underline{u})$ increases.*

The intuition for this result is that if $G(\underline{u})$ is low, few participants are at risk of quitting. By contrast, when $G(\underline{u})$ is high, workers and managers will quit unless they get their top choices. As a result, there are larger returns to arranging workers using their preferences when $G(\underline{u})$ is high. The base level of attrition $G(\underline{u})$ can vary for reasons with economic interpretations.

Corollary 2 (Job Quality)**.** *Let G' represent a distribution such that G first-order stochastically dominates G'. DA is more attractive to the CEO under G' than G.*

The shift in G' can be interpreted as making the firm less attractive to the workforce (draws from G' are more likely to lie below $\underline{u}$). If job quality is low, DA can offset this problem by placing workers into assignments they like. Selecting assignments with DA is a substitute for other job amenities; or conversely, being an attractive destination for workers is complementary with firm-dictated assignment.

Corollary 3 (Outside Options)**.** *DA is more attractive as the outside option $\underline{u}$ increases.*

Corollary 3 suggests that even firms with high job quality may nonetheless find DA attractive if the workforce's outside options are high enough.[20]

[20] One example of this is Google, a company that regularly appears at the top of Forbes' *Best Places to*

Corollary 4 (Asymmetric Information)**.** *Let G' be a mean-preserving spread of G, so that G' has the same mean of G but higher variance. Unless the base rate of quitting, $G(\underline{u})$, is above a critical threshold, DA is more attractive for G' than G.*

Corollary 4 addresses the amount of private information known to the workers, but not the CEO. A higher variance G corresponds to greater CEO uncertainty, and greater information asymmetry. Higher asymmetry increases the expected returns to DA.

Extension 2 (Replacement and Quitting Costs.) Finally, Appendix **??** shows that DA is more appealing in settings with high quitting or replacement costs (e.g., tight labor markets).

3.6 Alignment Between CEO & Workforce Preferences

Thus far, our framework has portrayed workforce preferences as private and independent of the CEO's (Assumption 4). We now relax this assumption by making the workforce aligned with the CEO – which could arise either through intrinsic motivation or external incentives. To formalize this type of alignment, we introduce two new ways of relating the workforce's preferences to the CEO's.

Definition 7 (Broad Alignment)**.** *An agent has* broadly aligned preferences with the firm *if their utility over matches is of the form $f(V) + \varepsilon_{ij}$ where $f(V)$ is an increasing function and ε_{ij} is uncorrelated with productivity. We say an agent has* strictly *broadly aligned preferences with the firm if their utility is of the form $f(V)$.*

"Broad aligned" workers gain utility from the company succeeding as a whole, and *strict* broad alignment means that the firm's overall success is the agent's *only* source of

Work list and offers workers free gourmet lunch and subsidized massages. Despite this, Google opened an internal talent marketplace in 2014 based on DA (Cowgill and Koning, 2018). Corollary 3 suggests this may be justified by the tight labor market for software engineers.

utility. Broad alignment would appear to be favorable for the use of DA, as it would give workers an incentive to use private information for the benefit of the firm.

However, broad alignment changes the technology of preferences. Broadly aligned workers have preferences over all $N!$ possible allocations; however, DA only permits workers to rank N members of the other side. As such, broadly aligned agents have interdependent values for their match partners. Worker i's payoff from matching with division j depends on how other workers and managers (beyond i and j) are matched. The group thus faces a coordination problem — even if they are highly aligned with the CEO.

Proposition 3. *If workers and divisions have strictly broadly aligned preferences, the CEO's optimal match is the surplus-maximizing Nash equilibrium of DA, but is not generally a unique Nash equilibrium. DA can equal the performance of firm-dictated assignment but cannot exceed it.*

The result highlights two aspects of delegating matches in organizations. First, when the workforce is broadly aligned, coordination problems arise through multiple equilibria.[21] Even when agents are maximally aligned, workers could fail to coordinate on the best equilibrium. Firm-dictated assignment may be preferable for coordination reasons.

Second, the benefit of delegating comes from tapping private information about match-specific quits. However, the workforce has the same preferences as the CEO under broad alignment, so there is no private information about preferences for DA to aggregate (and so no benefits). Broadly aligned agents are less likely to quit because the CEO's preferred match is also their most preferred match. Under strict broad alignment, the CEO can implement the optimal match without delegation (achieving the same outcome).

Broad alignment is a very strong assumption. In many other settings researchers have

[21] Proposition 3 examines *strict* broad alignment (an extreme case), but multiple equilibria issues could also arise if workers and divisions had private values for each match (and were paid a percentage of total firm output).

found strong firm-wide incentives difficult to create (Holmström, 1979; Oyer, 2004). In practice, a different form of alignment often appears.

Definition 8 (Narrow Alignment)**.** *An agent is said to have* narrowly aligned preferences with the firm *if their utility over matches is of the form* $f(v_{ij}) + \varepsilon_{ij}$ *where* $f(\cdot)$ *is an increasing function and* ε_{ij} *is uncorrelated with productivity. We say an agent has* strictly *narrowly aligned preferences with the firm if their utility is of the form* $f(v_{ij})$.

Narrowly aligned agents gain utility from their *own* productivity (irrespective of their co-workers' matches). This is similar to paying salespeople for their individual performance, and not for the success of the firm as a whole. Narrow alignment features private value preferences with some alignment with the CEO, but avoids the coordination issues of broad alignment (as agents do not have to consider the behavior of others). However, there are limitations to narrow alignment.

Proposition 4 (Hoarding)**.** *Narrow alignment is not sufficient or necessary to guarantee that DA yields output as high as FD.*

Even if all workers aspire to their most productive uses (individually), the result may underperform for the firm as a whole. The intuition for Proposition 4 is that narrow alignment encourages managers to hoard the most productive workers, even if these workers would have better use elsewhere in the company.[22] In the extreme case, narrow alignment on both sides of the match could produce an assortative match, which is optimal only for supermodular production functions (Becker, 1973).

In more general terms, narrow alignment is insufficient because of externalities. Each assignment imposes a *displacement externality* that affects other agents and the CEO. Even

[22]Our setting is motivated by horizontally differentiated placements, but managers could also "hoard" talent by denying promotions to deserving workers (Haegele, 2021). Friebel and Raith (2022) contains a model of talent hoarding. Similarly, narrow alignment by workers produces the reciprocal "career hoarding" (seeking out the most productive assignments, even if other workers would be more productive in these jobs).

if a worker performs a job well, their assignment may leave others in the company without productive uses. The CEO can internalize these externalities in firm-dictated assignment, but delegating to narrowly aligned agents does not automatically achieve this.

While aligned preferences are insufficient for delegation to perform well, narrowly aligned preferences are not *necessary* either. Appendix **??** shows an example where DA leads to an optimal configuration, despite worker preferences *not* being aligned according to our definition.[23] Despite these limitations, narrow alignment can be sufficient for DA to perform well for certain types of production technology. For example:

Corollary 5. *Suppose that workers and divisions have strictly narrowly aligned preferences, and that match productivities come from a production function that is supermodular in workers' and divisions' types. Then the worker-proposing DA selects the CEO-optimal assignment. However, the CEO can select this without DA using firm-dictated assignment.*

Corollary 5 shows a setting where delegation performs well. Unlike our broad alignment results, there is no coordination or multiple equilibria problem. However, like the broad alignment results, there is no tension between CEO and worker wants.

Proposition 3 and Corollary 5 show that DA and firm-dictated assignment sometimes produce the same outcome. In these settings, one may be preferable for reasons outside of our model. For example, setting up a market may involve fixed costs. Delegation may also have *benefits* outside the model. A longstanding view in organizational psychology and other disciplines is that workers have intrinsic, non-instrumental value for decision rights (Bartling et al., 2014). Anecdotes from Google suggest that workers value the symbolism of workers' choice, separately from its instrumental value (Cowgill and Koning, 2018). These issues are outside of our model but may be an important consideration in

[23] Of course, this raises the possibility that "negatively aligned" preferences could result in the optimal organizational match, even if preference-respecting mechanisms were used. In Appendix **??**, we provide an example that this is possible. In our example, workers and managers want to *avoid* working on the projects where they are individually most productive. Nonetheless, preference-respecting mechanisms produce the optimal assignments.

practice.

Which Side Proposes? Until now, all of our results – both with alignment and without – assume that both sides' preferences are symmetric. As such, the choice of the proposing side does not matter. However, it is possible that (e.g.) the manager's side is aligned with the CEO (broadly or narrowly), while the worker side's preferences are I.I.D. (as in Assumption 4). In this case, the proposing side could change outcomes from DA (although in many large markets, there is a unique stable match that DA selects no matter which side proposes).[24] In our empirical section, we compare outcomes from both workers-proposing and managers-proposing DA as a robustness check.

3.7 Correlated Preferences: Same-Side and Cross-Side

Finally, preferences could be private-value, but correlated within the workforce (rather than I.I.D.). We summarize intuition here and formalize in Appendix **??**. Proposition **??** shows that if one side has vertical preferences, then DA does not increase that side's retention. In Appendix **??**, we present simulation results about intermediate cases where the degree of same-sided correlation varies between zero and 1. Our results suggest that as preferences become highly correlated, DA's retention rate ultimately decreases to the unconditional retention rate.

Preferences can also be correlated across sides. If preferences are correlated across sides but are otherwise I.I.D. (i.e., no vertical preferences), then Equation 6 is the same, and the results in our I.I.D. model all hold. For additional details about these results, see Appendix **??**.

[24] For examples and evidence of uniqueness, see Mauras (2021); Ashlagi et al. (2017); Roth and Peranson (1999).

4 Empirical Setting

Our theoretical results highlight the potential tension between retention and match value for the firm. However, these tensions could be small or even non-existent in practice. If workers are aligned and coordinated, "dictating" and "delegating" could produce indistinguishable results. We now measure these forces in an applied setting. Is misalignment big enough to matter? If so, what drives the misalignment?

Our setting is a Fortune 500 company that develops software for business clients. The software includes customized tools for organizing and indexing special files and/or importing and managing content libraries. Employees are organized in teams serving a product and/or client. These teams feature a mixture of engineers and non-technical staff. Engineers hold a BS in computer science, and non-technical staff hold a BA in a social science, professional, or humanities subject. A single manager oversees each team. Prior to the adoption of an internal job marketplace, each participant was assigned to a team indefinitely.

Internal Mobility. During the sample period, the organization's leadership changed the policy of indefinite assignment. After employees spent several years on a team, they wanted career growth. Promotions increased pay and expectations, but left the underlying work unchanged (managers also sometimes wanted a new worker).

The company's leadership developed an internal talent market. All project assignments were given a pre-established "term length," measured in quarters. The average term length for positions on the market was one year. All workers whose term was ending — including those who were successful in their previous jobs — would go into the market to be reassigned.[25] Those who wished to stay in their previous roles could usually remain,

[25] "Term lengths" were created in part to avoid adverse selection in which only bad workers or managers sought new assignments. After assignments are made, workers have approximately 1-2 weeks to transition

but only if they were re-matched through the market.[26]

To avoid disruption, entry into the market was staggered. The firm aimed for no more than 25% of workers to be on the market at any time, so that the remaining 75% could focus on their day-to-day work. Because of staggering, each manager was typically recruiting at most one new worker in any given quarter's match, and the match was therefore one-to-one. Nearly all quarters' markets were unbalanced featuring an excess of managers.

To help participants make informed choices, the firm provided several sources of information. Each quarter, eligible participants developed a profile about their interests, accomplishments, and skills. Managers also included their job opening and skill requirements. Profiles could be searched or browsed through a web application. Although we could not obtain a screenshot for this paper without revealing internal information, Figure **??** includes a mockup that replicates a hypothetical user's profile page. In Appendix **??**, we document other sources of information that participants could use to make informed choices about prospective match partners.

The workers and managers were then matched using the worker-proposing deferred acceptance algorithm (Assumption 5). Each worker submitted a rank-ordered list of their preferred teams before a deadline, and the team managers submitted a similar ranking over workers. Figure **??** contains a replica of the submission page. The firm trained the workforce so that all sides could understand how their rankings would be used by the DA algorithm, emphasizing that all rankings would be kept strictly confidential. Although the firm broadly allowed workers on the market to transfer anywhere, the market was segmented into two submarkets per quarter based on specialization (Proposition 1) as either engineers or non-engineers.

Table 3 displays summary statistics of market participation. Our data consists of 318

between old and new assignments.

[26]In some cases, positions were ended after the term limit.

workers applying to 517 divisions/managers across seven submarkets. Participants ranked around 10% of options on average, leaving the remainder tied for last. The average worker was on the market 1.67 times, and all managers appeared only once.

Table 3: **Summary Statistics**

	% of choices ranked			# of times on the market			
	Mean	Min	Max	Mean	Min	Max	N
Workers	8.9	1.6	100	1.7	1	4	318
Managers	9.6	1.7	100	1	1	1	517

Notes: Summary statistics for our sample. Appendix **??** contain additional summary statistics.

4.1 Mapping the Setting to Theory

In our setting, workers and teams (led by the manager) are the two sides that need to be matched together 1:1 (as in our theory). Initial assignments were chosen by centralized administrators who performed a role similar to our theoretical CEO: While they hoped workers would like their assignments, they were uninformed about worker preferences, and optimized for the objectives of the company.

Our setting features several key characteristics appearing in our theoretical setup. For legal reasons, promotions and pay raises were handled separately from the market (Assumption 2, no match-specific pay). Workers and managers are relatively skilled and have good outside options (Corollary 3). The overall retention rate was high compared to industry standards (Proposition 2), but concerns about high replacement costs (Lemma **??**) led to the adoption of the jobs marketplace. Workers are paid in part with stock compensation to align incentives with shareholder goals (broad alignment, Def. 7). However, workers and managers are also paid a similar amount in individual bonuses based on their own performance (narrow alignment, Definition 8).[27]

[27] For the workers in our sample, approximately 40% of their annual compensation is performance based, with about half coming from stock options (broad alignment) and the other half coming from cash performance bonuses. This is a relatively high percentage of profit-sharing compared to the US labor force as a whole (Kurtulus and Kruse, 2017).

Our data include participant characteristics and preferences. Appendix **??** describes all variables in greater detail. The two sides' valuations for each other – μ^i_j and u^j_i cardinal utilities for match partners – are not directly observable in our setting. However the rankings submitted to DA could be interpreted as the ordering of each worker's μ^i_j (and each manager's u^j_i). During early rounds of the match, participants were surveyed anonymously about whether their rankings reflected their true preferences. 90% of managers confirmed that they did, and the remaining 10% reported feeling "neutral" about whether the reported preferences were true. Likewise, 72% of workers agreed that their preferences reflect their true preferences with 22% "neutral."

The firm's value for matching i and j — the v_{ij} terms — are critical elements of our model. Preferences of senior executives (v_{ij}) may diverge from those of workers (μ^i_j) and even of middle managers (u^j_i). In the next section, we present the empirical version of v_{ij} in our setting, and its connection to our model.

5 Empirical Strategy: Firm Match Quality

The goal of our empirical analysis is to quantify the differences between firm-dictated and market-generated matches, to measure the degree of misalignment, and to find reasons these matchings differ. A key input into these questions is the firm's measure of match quality, v_{ij} in our theory model (and particularly V_{FD} and V_{DA}, the outputs under firm-dictated assignment and DA). We now lay out our strategy for measuring these match-specific business values.

Research Design. Our empirical strategy is to evaluate matches using a match-specific productivity scoring function developed by the firm itself. The score was developed to evaluate match quality, and to encourage workers towards more preferred assignments (as we describe below). The firm's executives were worried about issues outlined in our theory framework: A match based only on workforce preferences could generate unproductive teams. For example, participants could optimize matching around socializing, rather than skill qualifications and performance. As a result, the firm's leaders directed its technical staff to develop a method for scoring the quality of any matching using methods similar to Graham (2011) and related papers (listed below).

We use the firm's own model to score match quality from the executive's perspective. A key benefit of this approach is that it was implemented by the firm itself. A classic idea in organizational economics is the need for firms to balance multiple, potentially competing objectives — some of which are hard to define. Rather than choosing these objectives (and their weights) ourselves as outside researchers, we leverage how the firm's leadership itself weighed these competing objectives. By using the firm's implementation, we leverage the local knowledge of the firm's leaders about their firm's business objectives and production function.

A drawback is that we cannot provide full diagnostics about the scoring function. However, we conduct a validation exercise below to study how the firm's scores map to realized outcomes. The scoring algorithm was developed by the firm and its engineering professionals using data and guidance we do not have access to. We disclose the relevant details about the score below, including how it was developed and used.

This approach is similar to a series of recent papers spanning multiple domains. In these papers, researchers estimate a match-specific productivity function from observational data. They then use the estimated function to simulate counterfactual assignments that address policy goals. Because of endogeneity in assignments in observational data, these authors develop innovative techniques (typically using natural experiments) to recover productivity functions (Graham 2011 reviews these methods). Recent applications include medical settings (Agarwal et al., 2020; Dahlstrand, 2021) and teacher assignment (Boyd et al., 2013; Aucejo et al., 2022).

Development of Match Quality Scores. The scoring algorithm was initiated and implemented by the firm itself, without our direct encouragement or involvement. The firm's match quality scores were based on the *characteristics* of workers and positions (such as their skill and experience profiles). In other settings, researchers have found that matching workers based on skills is a key benefit of management (Minni, 2022).

The firm used a structured skill and experience taxonomy to keep track of worker skills and job requirements. This is similar to the taxonomies in papers featuring O*NET data (Autor et al., 2003) or Burning Glass job postings (Deming and Kahn, 2018). Similar skill taxonomies are a common feature of HR software suites for enterprises (Minni, 2022).

The skill, background, and experience profile of each worker and manager – along with the requirements, job description and other structured attributes of each role – were used to predict a variety of outcomes, including sales revenue, engineering productivity metrics (e.g., lines of code contributed without bugs), customer satisfaction, response times, historical performance reviews, and other metrics of business performance. Because there are a variety of ways to weight these outcome metrics, the firm's researchers experimented with how to blend them, proceeding iteratively with feedback from the senior executives. The scoring function was developed from data, but also manually fine-tuned to incorporate difficult-to-quantify costs and benefits.

Analysts at the firm were aware that the endogeneity of historical matches could bias their estimates. To overcome these challenges, they used econometric methods for estimating match production functions using only data on observed matches, similar to Graham (2011). We did not help the firm estimate these models so we do not know the full set of estimation strategies or diagnostics they tested. We know that, historically, they have used sophisticated strategies, like leveraging natural experiments, for this type of estimation. For example, the firm had previously estimated match-specific productivity using a natural experiment where workers were assigned to projects based on their start dates. The mapping of start-dates to business units was not decided until a few days in advance of the employee's arrival (even by the firm itself). As such, two similar employees who arrived on slightly different dates could have very different project assignments.

Validating the firm objective score. In Appendix **??**, we examine realized outcomes, and we ask whether good outcomes for the firm were forecast by the objective function developed above. To study this, we use data about the required skills for each job, and how workers were subjectively assessed on the performance of each skill.

We find that a single standard deviation increase in the firm objective score corresponds to a 30 percentage point increase in the percentage of required skills assessed at performing or mastery level. This is an increase of 0.95 standard deviations. This (and other results in Appendix **??**) provides some evidence that the objective function was a reasonable proxy for positive outcomes for the firm.

As a robustness measure, we also develop simulations in Section 8 in which the firm objective score is *noisily* related to "true" match quality scores (i.e., an empirical analogue to Extension 1).

How the Firm Used the Scores. The firm used the scores to assess match quality in its new market. As in our theory, the firm was concerned with maximizing the sum of these scores across all realized assignments (Definitions 3 and 5). As such, the scores are an analogue to the v_{ij} terms in our model.

The leadership also deployed the scoring function inside the web application as color-coded icons. This was meant to encourage the workforce to rank productive matches favorably. When participants browsed the profiles of potential matches, the web application displayed an icon indicating the company's assessed match quality. Workers and middle-managers were told what these icons meant, so that workers who cared could follow the firm's suggestions. All potential matches were labeled, expressing the degree of company approval using color-coded icons. These are shown in the Figure **??** & **??** screenshot mockups.

The firm took the scores seriously enough to commit technical and executive resources developing them, to evaluate assignments with them, and to operationalize them as an encouragement inside the matching application. Importantly, we do not claim these are "correct" preferences for an executive to hold. Despite the validation results above, it is possible the executives valued matches that were (in some sense) imperfect. We claim only that these measures capture the executives' *preferences* over matches. Which set of preferences are "correct" — i.e., how the executives should prioritize multiple competing objectives — is a normative concept we do not take a stand on.

Effect of Icons. The presence of firm-sponsored guidance or encouragement is common in internal talent markets.[28] Workers in our setting were free to ignore the suggestions. Because these scores were shared with workers and managers, they might have influenced rankings in our setting. Insofar as the scores affected rankings, our results could be interpreted as a *lower bound* of the misalignment between workforce preferences and firm objectives.

Match Quality for the Firm. How is match quality distributed across potential assignments? Figure **??** shows $\mathcal{V}$ (Definition 2), the distribution of match quality scores. This graph suggests that the firm's overall performance is very sensitive to matchings (using the executives' preferred metric), and Appendix **??** quantifies the degree of match-specific productivity. These results suggest that the workforce is relatively specialized (needed to justify firm-dictated assignment under Proposition 1).

[28] For example, the internal marketplaces for clients of Gloat feature automated guidance for workers and managers as they select options.

Counterfactual Assignments. Using the match quality scores, we can now generate the set of potential firm-dictated assignments. To calculate the firm-dictated match, we use the Kuhn-Munkres algorithm that finds the matches that maximize the score. We include some other matching techniques for comparison. This includes a worker (or manager) draft (i.e., random serial dictatorship by one side), as well as fully random assignment with equal probabilities. We also examine a match quality minimizer, which identifies the matching with the lowest total firm objective score (while still ensuring all workers are matched). Each matching algorithm is run 50 times resolving ties randomly, generating a distribution of potential assignments. We bootstrap standard errors by market throughout our analysis. Appendix **??** describes all counterfactual assignment procedures in full detail.

6 Empirical Results

We now turn to the central question: How do the firm-dictated assignments compare to the self-organized matches using DA? Because the two methods were optimized for different outcomes, each will perform better on the outcomes they were tuned for. However, our theory model showed the key question is the magnitude of these differences (Equation 6). How much larger are the match quality benefits of centralized assignment (V_{FD}) versus delegation (V_{DA})? How much happier is the workforce under delegation versus centralized assignment (i.e., R_{DA} and R_{FD})? How large is the misalignment between the workforce and executives, and is it large enough to forgo the retention benefits of the market?

Table 4 compares the distribution of match quality scores obtained by deferred acceptance and firm-dictated assignment. The firm-dictated match has a 33% higher match quality score than random matches, while deferred acceptance is only about 5% higher.

Table 4: **Average Match Quality Score by Matching Algorithm**

	Mean	Min	P25	Median	P75	Max
Firm-Dictated	0.80	0.14	0.65	0.82	1.00	1.00
	(0.03)	(0.07)	(0.04)	(0.06)	(0.04)	(0.00)
Managers-Propose DA	0.63	0.00	0.43	0.65	0.82	1.00
	(0.05)	(0.01)	(0.04)	(0.07)	(0.11)	(0.00)
Workers-Propose DA	0.63	0.00	0.43	0.65	0.82	1.00
	(0.05)	(0.01)	(0.04)	(0.07)	(0.11)	(0.00)
Worker Draft	0.63	0.00	0.43	0.65	0.82	1.00
	(0.05)	(0.01)	(0.04)	(0.07)	(0.11)	(0.00)
Manager Draft	0.63	0.00	0.43	0.65	0.82	1.00
	(0.05)	(0.00)	(0.04)	(0.07)	(0.11)	(0.00)
Random Assignment	0.60	0.00	0.39	0.60	0.78	1.00
	(0.06)	(0.00)	(0.05)	(0.07)	(0.12)	(0.00)
Match Quality Minimizer	0.38	0.00	0.16	0.30	0.49	1.00
	(0.09)	(0.00)	(0.03)	(0.09)	(0.21)	(0.13)

Notes: This table displays the average match quality score for assignments generated by various algorithms. Standard errors were bootstrapped by the quarter of the market, using 1K bootstrapped samples per matching strategy. Appendix **??** contains variations of this table and a regression version.

Although we can reject the difference between DA and random matches, the size of the difference is small.

Although there is wide variability in each agents' best and worst match, we also find strong individual fixed effects on match value across (all possible match partners). We can use these individual fixed effects to measure the size of the DA vs FD difference, relative to the average quality of workers and managers. Rather than use the FD match, the firm could try to achieve the same increase in value by increasing the quality of hires. Our results indicate doing so would be difficult. FD matching results in a 27% improvement in the firm's match quality score over DA. Less than 20% of workers and 16% of managers have fixed effects larger than this amount. Increasing the average value of assignments by 26% is over 77% of a standard deviation in the worker fixed effects and 81% of a standard deviation in manager fixed effects, suggesting that achieving the same benefits through hiring or training would require a large increase in quality.

Workforce Preferences. If the firm were to dictate matches instead of using deferred acceptance, how would workers and managers fare? Tables 5 studies how workers ranked their assignments (and Table **??** for managers). Under DA, workers ranked their match 1.2 and managers ranked their match 8.2 on average.[29] Moreover, in DA matches, only 6% of workers and managers were assigned to a partner ranked as "tied for last."

Table 5: **Worker Ranking of Assignment, by Matching Algorithm**

	Mean	Min	P25	Median	P75	Max
Worker Draft	1.15	1.00	1.00	1.00	1.00	5.12
	(0.03)	(0.00)	(0.00)	(0.00)	(0.00)	(0.99)
Workers-Propose DA	1.17	1.00	1.00	1.00	1.00	8.00
	(0.04)	(0.00)	(0.00)	(0.00)	(0.00)	(0.96)
Managers-Propose DA	1.17	1.00	1.00	1.00	1.00	8.00
	(0.04)	(0.00)	(0.00)	(0.00)	(0.00)	(0.96)
Manager Draft	1.95	1.00	1.00	1.00	2.00	15.78
	(0.15)	(0.00)	(0.00)	(0.20)	(0.15)	(8.21)
Firm-Dictated	3.33	1.00	2.00	2.00	4.00	16.00
	(0.39)	(0.00)	(0.00)	(0.32)	(0.79)	(2.72)
Random Assignment	3.36	1.00	2.00	2.00	4.00	21.46
	(0.39)	(0.00)	(0.00)	(0.41)	(0.69)	(7.90)
Match Quality Minimizer	3.53	1.00	2.00	2.00	4.00	15.72
	(0.40)	(0.00)	(0.00)	(0.50)	(0.70)	(2.72)

Notes: This table displays how workers ranked their assignments under various algorithms. Standard errors were bootstrapped by the quarter of the market, using 1K bootstrapped samples per matching algorithm. All ties are assigned to the most favorable possible ranking. Appendix **??** contains variations of this table and a regression version. This table contains results for all managers, including those unmatched because of the excess of managers. We also created the table with matched managers only in Table **??**.

By contrast, 88% are in a tied for last match in the firm's preferred assignment. In percentile terms, switching from DA to the FD match means moving from the 3rd to the 44th percentile rank of choices for workers, and from the 17th to the 50th percentile rank of choices for managers (see Appendix **??**). Across measures, centrally planned matches are about as attractive to workers/managers as random assignments.

[29] The average ranking for managers is high because the table displays the average across all managers, including unmatched managers whose ranking is imputed as their last choice. If we only examine matched worker-manager pairs, the average manager ranking is 1.2 (see Table **??** in the appendix).

7 Why Do Assignments Differ?

Why does deferred acceptance produce such different matches than firm-dictated assignment? There could be many reasons. Below we examine three: misalignment between executive and workforce preferences, assortative matching, and the desire for skill growth through on-the-job training.

7.1 Misalignment

Matches could differ because of how workers are "aligned" with the firm's objectives, either in the narrow (Definition 8) or broad sense (Definition 7). In Table 6, we run OLS specifications predicting which option is each agent's #1 choice. The data in these regressions include all possible assignments for each worker and manager.

We focus on two key explanatory variables. The first is the firm's match quality score (corresponding to the v_{ij} in the model). We standardize this term to aid interpretation. The resulting coefficient suggests how much each worker is narrowly aligned. Second, we include a variable measuring broad alignment for each possible i, j match (also standardized). The measure of broad alignment is whether each i, j pair is present in the firm-dictated assignments that maximize overall match quality (using Kuhn-Munkres).[30]

Table 6 shows results for both types of alignment. Across both measures and specifications, the coefficients are economically small in magnitude and do not explain a high fraction of variance in preferences. In Table **??**, we find similar results using rank-ordered logistic specifications. Although we can statistically reject zero, our point estimates and standard errors rule out large relationships across our results. For these reasons, the workforce does value matches with a higher firm objective score, but executive preferences do

[30] In some cases, there are multiple ways to achieve a maximum. In these cases, we express "broad alignment" of a match i, j as the proportion of all optimal matches that include i, j.

not appear to be a major driver of workforce preferences.

Misalignment or Coordination Failure? In our theoretical framework, a broadly aligned workforce could settle on a bad equilibrium. In Appendix **??**, we examine whether workers and managers make firm-optimal choices, conditional on what other players are doing. We find little evidence that they do. In Appendix **??**, we find that there are, in fact, a set of beliefs about other players' actions that would justify the observed behavior as better-than-random for the firm. However, these beliefs are no more accurate than random uniform draws from the set of all possible preference profiles.[31] These results provide some suggestive evidence that the observed misalignment is not driven by multiple equilibria and coordination failures around common goals.

Table 6: **Preference Alignment**

	Worker ranked manager #1			Manager ranked worker #1		
Objective Score (v_{ij}), σ	0.0024***		0.0013**	0.0028***		0.0017**
	(0.00063)		(0.00061)	(0.00071)		(0.00070)
Firm-Dictated Assignment, σ		0.0091***	0.0089***		0.010***	0.0098***
		(0.0021)	(0.0021)		(0.0022)	(0.0022)
Observations	23361	23361	23361	23361	23361	23361
R^2	0.000	0.004	0.004	0.000	0.004	0.005

Notes: This table examines the alignment between worker/manager preferences and the firm's preferences. Both explanatory variables have been standardized. The units of observation are a worker$\times$ manager pair. Columns 1–3 display the results of an OLS regression of whether the worker ranked the manager #1 as a function of the firm's match quality score (column 1), their match in the firm's preferred match (column 2), and both (column 3). Columns 4–6 display similar results for whether the manager ranked the worker #1. Table **??** shows similar results using rank-ordered logit specifications. Robust standard errors clustered at the worker level in columns 1–3, while robust standard errors clustered at the manager level in columns 4–6.

*** $p < 0.01$ ** $p < 0.05$ * $p < 0.10$

7.2 Assortative Matching

In Table 7, we directly assess the levels of assortativeness across different matchings. To measure the quality of each participant, we use their fixed effect on match quality scores (averaging across all possible partners). We find that the firm-dictated solution is more likely to assign high quality workers with lower quality managers, and vice versa. This suggests that the match quality scores are submodular. When we test the scores directly for submodularity in Appendix **??** we find they are indeed submodular. We can now see one of the reasons deferred acceptance differed: it produced a positive assortative match.

This is likely the result of participants being slightly "narrowly aligned" (Definition 7), as measured in our earlier results (Table 6). Proposition 4 shows that narrow alignment

[31] In addition, these beliefs are not the same for all members. Each participant would have their own, distinct inaccurate beliefs that are uncorrelated with other players' distinctive (also inaccurate) beliefs.

Table 7: **Assortative Matching by Algorithm**

	Mean	Min	P25	Median	P75	Max
Match Quality Minimizer	0.71	0.00	0.18	0.59	1.10	2.44
	(0.07)	(0.00)	(0.08)	(0.12)	(0.09)	(0.05)
Random Assignment	0.69	0.00	0.21	0.57	1.03	3.08
	(0.07)	(0.00)	(0.07)	(0.12)	(0.09)	(0.06)
Firm-Dictated	0.67	0.01	0.17	0.46	0.89	3.85
	(0.09)	(0.00)	(0.05)	(0.12)	(0.12)	(0.24)
Manager Draft	0.66	0.00	0.19	0.51	0.97	2.97
	(0.07)	(0.00)	(0.07)	(0.11)	(0.11)	(0.13)
Workers-Propose DA	0.63	0.00	0.18	0.50	0.93	2.66
	(0.06)	(0.00)	(0.06)	(0.11)	(0.09)	(0.29)
Managers-Propose DA	0.63	0.00	0.18	0.50	0.93	2.66
	(0.06)	(0.00)	(0.06)	(0.11)	(0.09)	(0.29)
Worker Draft	0.63	0.00	0.18	0.47	0.93	2.69
	(0.06)	(0.00)	(0.06)	(0.11)	(0.10)	(0.28)

Notes: This table examines the average absolute value of the difference between worker and manager quality generated. To measure the quality of each participant, we use their fixed effect on match quality scores (averaging across all possible partners). Algorithms are sorted from most to least assortative. Standard errors were bootstrapped by the quarter of the market, using 1K bootstrapped samples per matching strategy.

generates talent hoarding: all participants want a high-value partner, even if the CEO prefers talent to be spread out (because of submodular production technology). Together, these results show that assortative matching is a key reason why delegation failed to produce higher output in this setting.

7.3 Growth and Firm-Sponsored Training

What does the workforce value, if not the principal's goals? Our final result suggests that workers prioritized roles that create opportunities for on-the-job practice of key skills they lack, but desire to improve. Our evidence comes from two sources: 1) an anonymous survey administered to workers, and 2) data from workers' profiles where they could label which skills they would like to grow. We study this question in detail in Appendix **??**, and summarize our findings here for brevity.

Workers are more likely to designate a skill for growth if they do not already perform it well (Table **??**). In Table 8, we show that workers rank jobs higher that involve one of their growth skills, even though these jobs require skills they do not currently perform well. When we manually categorized each skill as general or firm-specific, we found that a vast majority of skills workers sought for growth (≈85%) were *not* firm specific.[32] As such, the workers sought to use the market to develop skills they could possibly market to other employers (even if it were not their initial intention to leave). By contrast, the firm's match quality score prefers workers use the skills where they already perform well, without the need for on-the-job training.

Our results show that the firm did not prioritize on-the-job-training in its match quality score. We do not take a stand on whether this was a "mistake." However, this finding is consistent with prior research about employer-sponsored training for non-firm specific

[32]That is, most skills workers desired to grow were general skills such as programming in a specific language such as Java or Python, or general business tasks such as spreadsheet analytics.

skills. The idea that employer-sponsored training is underprovided is a classic human capital topic (Pigou, 1912; Becker, 1964), particularly for non-firm specific skills.

Table 8: **Worker Preferences for Growth Jobs**

	Worker i Ranking of Manager j	Worker i Ranked Manager j #1
Number of i's Growth Skills featured in j's Job Description, σ	0.013***	0.0049**
	(0.0037)	(0.0020)
R^2	0.000	0.021
Observations	23361	23361
Estimation Approach	Ranked Logit	OLS

Notes: This table examines whether workers prioritize assignments that use skills the worker desires to grow. The units of observation are a worker× manager pair. All regressions control for the degree of broad and narrow alignment (as specified in Table 6 and Section 7.1). The OLS specification includes worker and manager fixed effects. Robust standard errors clustered by worker are included. For additional details, see Appendix **??**.

*** $p < 0.01$ ** $p < 0.05$ * $p < 0.10$

8 Extension: Noisily-Informed Executives

How much do our empirical results depend on executives being well-informed? To answer this, we perform simulations in Appendix **??** in which the firm's match quality score is a noisy proxy for value (modeled theoretically in Extension 1). This simulates the idea that executives optimized for the "wrong" outcome, with varying degrees of correlation from the "true" measure of productivity (in this scenario, our earlier results are an upper-bound). We show that if the correlation between the company's match quality score signal and "true" match quality is at least $\rho = 0.3$, we can still statistically reject the difference between DA and firm-dictated assignment in our setting. While this does not prove that firm-dictated assignment is "better," it does provide some quantification of how much precision is needed. The benefits of firm-dictated assignment appear to be robust to noise (in our setting). However, as measurement error increases, the planner's value of dictating diminishes to zero.

9 Discussion

Organizations often form teams by matching agents. Leaders may have goals for the organization as a whole. However, worker and managers' satisfaction are important constraints. This paper has offered a theoretical framework about a CEO's choice of assignment mechanisms and an empirical case study of the main tradeoff in the model.

The key trade-off we study is the balance between match value and job satisfaction (particularly retention). Our model connects the adoption of internal markets to firm characteristics including match-specific productivity (specialization), incentive alignment, information asymmetry, and production technology. Overall, our results point to new opportunities for market designers to work on organizational applications.

In our empirical results, we find a high degree of match-specific productivity and specialization. As a result, there are large potential gains in match quality from the executive's perspective. However, workers and managers are apathetic about these assignments. Our results suggest that these differences arise in part through differences in assortative matching. In the workforce driven match, the firm's best workers and managers team up together. However from the CEO's perspective, a good manager is more helpful in carrying the bad workers. We also find workers prioritize opportunities for on-the-job skill development – especially in non-firm specific skills – and the firm does not.

Our paper has implications for how trends in labor markets and technology affect internal organization and principal-agent problems inside organizations. A literature about job design and information technology documents a secular rise in generalist job design and multitasking,[33] partly driven by the automation of routine work (RBTC Autor et al.,

[33] Osterman (1994); Caroli and Van Reenen (2001); Deming (2017). As Dessein and Santos (2006) note, the biggest management fad of the 1990s, reengineering (Hammer and Champy, 2009), prescribes "combining several jobs into one" and thus "putting back together again the work that Adam Smith and Henry Ford broke into tiny pieces."

2003). Other researchers suggest secular trends towards higher quit rates (Fuller and Kerr, 2022). Both these trends are linked to increased adoption of internal markets in our model, suggesting that organizations will continue to explore market-like mechanisms for assignment problems.

Team formation and matching problems are abundant in organizations. Modern workers often have strong preferences about their job assignments; they care about career growth, interpersonal compatibility with supervisors and peers, and alignment with their social and political values. Our paper hopes to contribute a better understanding of how firms can use markets (or other institutions) to integrate workers' preferences with organizational priorities.

References

Abdulkadiroglu, Atila, Umut M Dur, and Aram Grigoryan, "School Assignment by Match Quality," Working Paper 28512, NBER February 2021.

Adhvaryu, Achyuta, Vittorio Bassi, Anant Nyshadham, and Jorge A Tamayo, "No line left behind," Technical Report, National Bureau of Economic Research 2020.

Agarwal, Nikhil, Charles Hodgson, and Paulo Somaini, "Choices and outcomes in assignment mechanisms," Technical Report, National Bureau of Economic Research 2020.

Allocca, Alessandra, "'No Man is an Island': An Empirical Study on Team Formation and Performance," *Available at SSRN 4365260*, 2023.

Alonso, Ricardo and Niko Matouschek, "Optimal delegation," *The Review of Economic Studies*, 2008, *75* (1), 259–293.

Arya, Anil, Jonathan Glover, and Bryan R Routledge, "Project assignment rights and incentives for eliciting ideas," *Management Science*, 2002, *48* (7), 886–899.

Ashlagi, Itai, Yash Kanoria, and Jacob D Leshno, "Unbalanced random matching markets: The stark effect of competition," *Journal of Political Economy*, 2017, *125* (1), 69–98.

Aucejo, Esteban, Patrick Coate, Jane Cooley Fruehwirth, Sean Kelly, and Zachary Mozenter, "Teacher effectiveness and classroom composition: Understanding match effects in the classroom," *The Economic Journal*, 2022, *132* (648), 3047–3064.

Autor, David H, Frank Levy, and Richard J Murnane, "The skill content of recent technological change," *The Quarterly journal of economics*, 2003, *118* (4), 1279–1333.

Ba, Bocar, Patrick Bayer, Nayoung Rim, Roman Rivera, and Modibo Sidibé, "Police Officer Assignment and Neighborhood Crime," Technical Report, NBER 2021.

Baker, George and Bengt Holmstrom, "Internal labor markets: Too many theories, too few facts," *The American Economic Review*, 1995, *85* (2), 255–259.

Bandiera, Oriana, Iwan Barankay, and Imran Rasul, "Social Incentives in the Workplace," *The Review of Economic Studies*, 2010, *77* (2), 417–458.

Barron, Greg and Felix Vardy, "The Internal Job Market of the IMF's Economist Program," *IMF Staff Papers*, 2005, *52* (3), 410–429.

Bartling, Björn, Ernst Fehr, and Holger Herz, "The Intrinsic Value of Decision Rights," *Econometrica*, 2014, *82* (6), 2005–2039.

Bates, Michael D, Michael Dinerstein, Andrew C Johnston, and Isaac Sorkin, "Teacher Labor Market Equilibrium and Student Achievement," Technical Report, NBER 2022.

Becker, G, "Human Capital, Chicago: The University of Chicago Press," *Journal of Human Capital*, 1964, *9* (3).

Becker, Gary S, "A Theory of Marriage: Part I," *Journal of Political Economy*, 1973, *81* (4), 813–846.

Bellet, Clement S, Jan-Emmanuel De Neve, and George Ward, "Does employee happiness have an impact on productivity?," *Management science*, 2023.

Benson, Alan, Danielle Li, and Kelly Shue, "Promotions and the peter principle," *The Quarterly Journal of Economics*, 2019, *134* (4), 2085–2134.

_ , _ , and _ , "Potential and the gender promotion gap," *Unpublished Working Paper*, 2021.

Bloesch, Justin, Birthe Larsen, and Bledi Taska, "Which workers earn more at productive firms? Position specific skills and individual worker hold-up power," *Position Specific Skills and Individual Worker Hold-up Power (March 1, 2022)*, 2022.

Bloom, Nicholas and John Van Reenen, "Human resource management and productivity," in "Handbook of labor economics," Vol. 4, Elsevier, 2011, pp. 1697–1767.

Boyd, Donald, Hamilton Lankford, Susanna Loeb, and James Wyckoff, "Analyzing the Determinants of the Matching of Public School Teachers to Jobs : Disentangling the Preferences of Teachers and Employers," *Journal of Labor Economics*, 2013, *31* (1), 83–117.

Bresnahan, Timothy F, Erik Brynjolfsson, and Lorin M Hitt, "Information technology, workplace organization, and the demand for skilled labor: Firm-level evidence," *The Quarterly Journal of Economics*, 2002, *117* (1), 339–376.

Brynjolfsson, Erik and Lorin M Hitt, "Beyond computation: Information technology, organizational transformation and business performance," *Journal of Economic perspectives*, 2000, *14* (4), 23–48.

Cappelli, Peter, "HR for neophytes," *Harvard Business Review*, 2013, *91* (10), 25–27.

Caroli, Eve and John Van Reenen, "Skill-biased organizational change? Evidence from a panel of British and French establishments," *The Quarterly Journal of Economics*, 2001, *116* (4), 1449–1492.

Cassar, Lea and Stephan Meier, "Nonmonetary incentives and the implications of work as a source of meaning," *Journal of Economic Perspectives*, 2018, *32* (3), 215–38.

Cestone, Giacinta, Chiara Fumagalli, Francis Kramarz, Giovanni Pica et al., "Exploiting Growth Opportunities: The Role of Internal Labor Markets," *The Review of Economic Studies*, 2022.

Christensen, Michael and Thorbjørn Knudsen, "Design of decision-making organizations," *Management Science*, 2010, *56* (1), 71–89.

Combe, Julien, Umut Mert Dur, Olivier Tercieux, Camille Terrier, M Utku Ünver et al., "Market Design for Distributional Objectives in (Re) assignment," Technical Report, Boston College Department of Economics 2022.

Cowgill, Bo and Rembrand Koning, "Matching Markets for Googlers," March 2018. Harvard Business School Case 718-487.

_, Zoë Cullen, Anh Nguyen, and Ricardo Perez-Truglia, "Gender Segregation in the Workplace," *Working paper*, 2020.

Dahlstrand, Amanda, "Defying Distance? The Provision of Services in the Digital Age," 2021.

Davis, Jonathan M.V., "Labor Market Design Can Improve Match Outcomes: Evidence From Matching Teach For America Teachers to Schools," 2022.

_, Kyle Greenberg, and Damon Jones, "An Experimental Evaluation of Deferred Acceptance: Evidence from Over 100 Army Officer Labor Markets," *NBER Working Paper*, 2023, *31612.*

Delgado, William, "Stable Assignment Problems for Staffing: Negotiated Complete Stable Matchings," *Chapter 2 of PhD Dissertation (Mallory Soldner, Georgia Tech)*, 2014.

Deming, David and Lisa B Kahn, "Skill requirements across firms and labor markets: Evidence from job postings for professionals," *Journal of Labor Economics*, 2018, *36* (S1), S337–S369.

Deming, David J, "The growing importance of social skills in the labor market," *The Quarterly Journal of Economics*, 2017, *132* (4), 1593–1640.

Dessein, Wouter, "Authority and communication in organizations," *The Review of Economic Studies*, 2002, *69* (4), 811–838.

— **and Tano Santos**, "Adaptive organizations," *Journal of political Economy*, 2006, *114* (5), 956–995.

Friebel, Guido and Michael Raith, "Talent Management: the Role of Bosses," Technical Report, Working Paper 2022.

Fuller, Joseph and William Kerr, "The Great Resignation Didn't Start with the Pandemic," *Harvard business review*, 2022.

Gale, David and Lloyd S . Shapley, "College Admissions and the Stability of Marriage," *The American Mathematical Monthly*, 1962, *69* (1), 9–15.

Graham, Bryan S, "Econometric methods for the analysis of assignment problems in the presence of complementarity and social spillovers," *Handbook of social economics*, 2011, *1*, 965–1052.

Haegele, Ingrid, "Talent Hoarding in Organizations," 2021.

Hamilton, Barton H, Jack A Nickerson, and Hideo Owan, "Team incentives and worker heterogeneity," *Journal of Political Economy*, 2003, *111* (3), 465–497.

Hammer, Michael and James Champy, *Reengineering the Corporation: Manifesto for Business Revolution, A*, Zondervan, 2009.

Holmström, Bengt, "Moral hazard and observability," *The Bell journal of economics*, 1979, pp. 74–91.

Hsieh, Chang-Tai and Peter J Klenow, "Misallocation and manufacturing TFP in China and India," *The Quarterly journal of economics*, 2009, *124* (4), 1403–1448.

Immorlica, Nicole and Mohammad Mahdian, "Marriage, Honesty, and Stability," in "Proceedings of the Sixteenth Annual ACM-SIAM Symposium on Discrete Algorithms" SODA '05 Society for Industrial and Applied Mathematics USA 2005, p. 53–62.

Impink, Stephen Michael, Andrea Prat, and Raffaella Sadun, "Communication within firms: evidence from CEO turnovers," *National Bureau of Economic Research*, 2021, *29042.*

Knuth, Donald Ervin, *Mariages stables et leurs relations avec d'autres problemes combinatoires*, Les Presses de l'Université de Montréal, 1976.

Kojima, Fuhito and Parag A Pathak, "Incentives and stability in large two-sided matching markets," *American Economic Review*, 2009, *99* (3), 608–27.

— **, Ning Sun, and Ning Neil Yu**, "Job matching under constraints," *American Economic Review*, 2020, *110* (9), 2935–47.

Kuhn, Harold W, "The Hungarian method for the assignment problem," *Naval research logistics quarterly*, 1955, *2* (1-2), 83–97.

Kurtulus, Fidan Ana and Douglas L Kruse, *How did employee ownership firms weather the last two recessions?: Employee ownership, employment stability, and firm survival: 1999-2011*, WE Upjohn Institute, 2017.

Lazear, Edward P, Kathryn L Shaw, and Christopher T Stanton, "The value of bosses," *Journal of Labor Economics*, 2015, *33* (4), 823–861.

Lee, Michael Y and Amy C Edmondson, "Self-managing organizations: Exploring the limits of less-hierarchical organizing," *Research in organizational behavior*, 2017, *37*, 35–58.

Lindbeck, Assar and Dennis J Snower, "Multitask learning and the reorganization of work: From tayloristic to holistic organization," *Journal of labor economics*, 2000, *18* (3), 353–376.

Martela, Frank, "What makes self-managing organizations novel?," *Journal of Organization Design*, 2019, *8* (1), 1–23.

Mas, Alexandre and Enrico Moretti, "Peers at work," *American Economic Review*, 2009, *99* (1), 112–145.

Mauras, Simon, "Two-sided random matching markets: Ex-ante equivalence of the deferred acceptance procedures," *ACM Transactions on Economics and Computation*, 2021, *9* (4), 1–14.

Minni, V, "Making the Invisible Hand Visible: Managers and the Allocation of Workers to Jobs," Technical Report, mimeo 2022.

Ortega, Jaime, "Job rotation as a learning mechanism," *Management Science*, 2001, *47* (10), 1361–1370.

Osterman, Paul, "How common is workplace transformation and who adopts it?," *Ilr Review*, 1994, *47* (2), 173–188.

Oswald, Andrew J, Eugenio Proto, and Daniel Sgroi, "Happiness and productivity," *Journal of labor economics*, 2015, *33* (4), 789–822.

Oyer, Paul, "Why do firms use incentives that have no incentive effects?," *The Journal of Finance*, 2004, *59* (4), 1619–1650.

Park, Sangyoon, "Socializing at work: Evidence from a field experiment with manufacturing workers," *American Economic Journal: Applied Economics*, 2019, *11* (3), 424–55.

Perry, Elissa L and Carol T Kulik, "The devolution of HR to the line," *The International Journal of Human Resource Management*,, 2008, *19* (2), 262–273.

Pigou, Arthur Cecil, *Wealth and welfare*, Macmillan and Company, limited, 1912.

Pittel, Boris, "On Likely Solutions of a Stable Marriage Problem," *The Annals of Applied Probability*, 1992, *May*, 358–401.

Raveendran, Marlo, Phanish Puranam, and Massimo Warglien, "Division of labor through self-selection," *Organization Science*, 2021.

Roth, Alvin E., "Deferred acceptance algorithms: history, theory, practice, and open questions," *International Journal of Game Theory*, jan 2008, *36* (3-4), 537–569.

Roth, Alvin E and Elliott Peranson, "The Redesign of the Matching Market for American Physicians: Some Engineering Aspects of Economic Design," *The American Economic Review*, 1999, *89* (4), 748–780.

Roth, Alvin E. and Marilda Sotomayor, *Two-Sided Matching: A Study in Game-Theoretic Modeling and Analysis*, Cambridge, UK: Cambridge University Press, 1990.

Sonmez, Tayfun and Tobias B Switzer, "Matching With (Branch-of-Choice) Contracts at the United States Military Academy," *Econometrica*, 2013, *81* (2), 451–488.

Thakur, Ashutosh, "Matching in the civil service," Technical Report, ECONtribute Discussion Paper 2021.

Whittaker, Susan and Mick Marchington, "Devolving HR responsibility to the line: threat, opportunity or partnership?," *Employee Relations*, 2003.

Xu, Guo, Marianne Bertrand, and Robin Burgess, "Organization of the State: Home Assignment and Bureaucrat Performance," *The Journal of Law, Economics, and Organization*, 10 2021.