

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

Working Paper

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

CESifo Working Paper, No. 11120

Provided in Cooperation with:

Ifo Institute – Leibniz Institute for Economic Research at the University of Munich

Suggested Citation: Cowgill, Bo; Davis, Jonathan M. V.; Montagnes, B. Pablo; Perkowski, Patryk (2024) : Stable Matching on the Job? Theory and Evidence on Internal Talent Markets, CESifo Working Paper, No. 11120, CESifo GmbH, Munich

This Version is available at:

<https://hdl.handle.net/10419/300048>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

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

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

Impressum:

CESifo Working Papers

ISSN 2364-1428 (electronic version)

Publisher and distributor: Munich Society for the Promotion of Economic Research - CESifo GmbH

The international platform of Ludwigs-Maximilians University's Center for Economic Studies and the ifo Institute

Poschingerstr. 5, 81679 Munich, Germany

Telephone +49 (0)89 2180-2740, Telefax +49 (0)89 2180-17845, email office@cesifo.de

Editor: Clemens Fuest

<https://www.cesifo.org/en/wp>

An electronic version of the paper may be downloaded

- from the SSRN website: www.SSRN.com
- from the RePEc website: www.RePEc.org
- from the CESifo website: <https://www.cesifo.org/en/wp>

Stable Matching on the Job?

Theory and Evidence on Internal Talent Markets

Abstract

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-Codes: M500, D470, J400.

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

Bo Cowgill
Columbia Business School
New York City / NY / USA
bo@bocowgill.org

Jonathan M. V. Davis
University of Oregon, Eugene / OR / USA
jdavis5@uoregon.edu

B. Pablo Montagnes
Emory University, Atlanta / GA / USA
pablo.montagnes@emory.edu

Patryk Perkowski
Yeshiva University, New York City / NY / USA
patryk.perkowski@yu.edu

May 15, 2024 (First Version: February 2021)

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](#); [Haeghele, 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

¹We 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.

²This 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

³These 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.

⁴Gloat'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,

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

⁶For readers unfamiliar with the deferred acceptance algorithm (Gale and Shapley, 1962; Roth and Sotomayor, 1990), Appendix A.2 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,

⁷The 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. \(2020\)](#).

⁸For 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 ([Haeghele, 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-

⁹Self-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-

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

¹¹Other 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 A.

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 $i \in \{1, \dots, I\}$ and J divisions $j \in \{1, \dots, J\}$. 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).

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

¹³For 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 :*

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

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 (2×2 Example)

		Division 1	Division 2
Worker 1	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 2×2 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).

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

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

Table 1: CEO Payoffs (2×2 Example)

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

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.

¹⁴In 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.

¹⁵The 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).

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

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

Definition 5 (Expected CEO Payoffs).

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

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:

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

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.

¹⁶For 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_i^j 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 A.9, 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 A.10, we also allow outside offers to be correlated with productivity.¹⁷

Structure of Assignment Mechanisms. To implement “firm-dictated assignment,” we show in Appendix A.1 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 A.2 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.

¹⁸See 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:

$$\frac{V_{FD}}{\bar{V}} \geq \frac{R_{DA}}{\bar{R}}. \quad (6)$$

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 A.4, 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 A1 shows that noisy measurement diminishes the benefits of firm-dictated

¹⁹In 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 A1 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

²⁰One 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 A.5 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

²¹Proposition 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

²²Our setting is motivated by horizontally differentiated placements, but managers could also “hoard” talent by denying promotions to deserving workers (Haeghele, 2021). Friebe 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 A.6 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

²³Of 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 A.7, 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 A.8. Proposition A2 shows that if one side has vertical preferences, then DA does not increase that side’s retention. In Appendix A.8, 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 A.8.

²⁴For 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 B1 includes a mockup that replicates a hypothetical user's profile page. In Appendix B.2, 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 B2 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.

²⁶In 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			N
	Mean	Min	Max	Mean	Min	Max	
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 C.2 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 A2) 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

²⁷For 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 B.3 describes all variables in greater detail. The two sides’ valuations for each other – μ_j^i and u_i^j 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_j^i (and each manager’s u_i^j). 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_j^i) and even of middle managers (u_i^j). 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 C.1, 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 C.1) 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 B1 & B2

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 C5 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 C.11 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).

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

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

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 [B.3](#) 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. Although we can reject the difference between DA and random matches, the size of the difference is small.

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

	Mean	Min	P25	Median	P75	Max
Firm-Dictated	0.80 (0.03)	0.14 (0.07)	0.65 (0.04)	0.82 (0.06)	1.00 (0.04)	1.00 (0.00)
Managers-Propose DA	0.63 (0.05)	0.00 (0.01)	0.43 (0.04)	0.65 (0.07)	0.82 (0.11)	1.00 (0.00)
Workers-Propose DA	0.63 (0.05)	0.00 (0.01)	0.43 (0.04)	0.65 (0.07)	0.82 (0.11)	1.00 (0.00)
Worker Draft	0.63 (0.05)	0.00 (0.01)	0.43 (0.04)	0.65 (0.07)	0.82 (0.11)	1.00 (0.00)
Manager Draft	0.63 (0.05)	0.00 (0.00)	0.43 (0.04)	0.65 (0.07)	0.82 (0.11)	1.00 (0.00)
Random Assignment	0.60 (0.06)	0.00 (0.00)	0.39 (0.05)	0.60 (0.07)	0.78 (0.12)	1.00 (0.00)
Match Quality Minimizer	0.38 (0.09)	0.00 (0.00)	0.16 (0.03)	0.30 (0.09)	0.49 (0.21)	1.00 (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 C.7 contains variations of this table and a regression version.

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 C10 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 (0.03)	1.00 (0.00)	1.00 (0.00)	1.00 (0.00)	1.00 (0.00)	5.12 (0.99)
Workers-Propose DA	1.17 (0.04)	1.00 (0.00)	1.00 (0.00)	1.00 (0.00)	1.00 (0.00)	8.00 (0.96)
Managers-Propose DA	1.17 (0.04)	1.00 (0.00)	1.00 (0.00)	1.00 (0.00)	1.00 (0.00)	8.00 (0.96)
Manager Draft	1.95 (0.15)	1.00 (0.00)	1.00 (0.00)	1.00 (0.20)	2.00 (0.15)	15.78 (8.21)
Firm-Dictated	3.33 (0.39)	1.00 (0.00)	2.00 (0.00)	2.00 (0.32)	4.00 (0.79)	16.00 (2.72)
Random Assignment	3.36 (0.39)	1.00 (0.00)	2.00 (0.00)	2.00 (0.41)	4.00 (0.69)	21.46 (7.90)
Match Quality Minimizer	3.53 (0.40)	1.00 (0.00)	2.00 (0.00)	2.00 (0.50)	4.00 (0.70)	15.72 (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 C.8 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 C11.

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 C.8.2). Across measures, centrally planned matches are about as attractive to workers/managers as random assignments.

²⁹The 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 C11 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 C16, 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

³⁰In 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 C.3, 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 C.4, 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*** (0.00063)		0.0013** (0.00061)	0.0028*** (0.00071)		0.0017** (0.00070)
Firm-Dictated Assignment, σ		0.0091*** (0.0021)	0.0089*** (0.0021)		0.010*** (0.0022)	0.0098*** (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 C16 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 C.6 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

³¹In 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 (0.07)	0.00 (0.00)	0.18 (0.08)	0.59 (0.12)	1.10 (0.09)	2.44 (0.05)
Random Assignment	0.69 (0.07)	0.00 (0.00)	0.21 (0.07)	0.57 (0.12)	1.03 (0.09)	3.08 (0.06)
Firm-Dictated	0.67 (0.09)	0.01 (0.00)	0.17 (0.05)	0.46 (0.12)	0.89 (0.12)	3.85 (0.24)
Manager Draft	0.66 (0.07)	0.00 (0.00)	0.19 (0.07)	0.51 (0.11)	0.97 (0.11)	2.97 (0.13)
Workers-Propose DA	0.63 (0.06)	0.00 (0.00)	0.18 (0.06)	0.50 (0.11)	0.93 (0.09)	2.66 (0.29)
Managers-Propose DA	0.63 (0.06)	0.00 (0.00)	0.18 (0.06)	0.50 (0.11)	0.93 (0.09)	2.66 (0.29)
Worker Draft	0.63 (0.06)	0.00 (0.00)	0.18 (0.06)	0.47 (0.11)	0.93 (0.10)	2.69 (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 C.13, 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 C18). 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 ($\approx 85\%$) 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

³²That 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*** (0.0037)	0.0049** (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 \times 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 C.13.

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

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 C.12 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.
- Haeghele, 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.

Appendix: For Online Publication Only

A Theory Appendix

A.1 Formalization of Firm-Dictated Assignment using the [Kuhn-Munkres Algorithm](#)

Any proposed match between worker i and division j has probability $P(i \leftrightarrow j \text{ acceptable})$ of being acceptable to both the worker and division, and payoff of $v_{i,j}$, for an expected value of $v_{ij}P(i \rightarrow j \text{ acceptable})$. Because the principal's objective in Equation 5 is to find matches that maximize the sum of these, the problem is equivalent to the assignment problem ([Kuhn, 1955](#)), even after integrating beliefs about quitting. It is thus solvable through the [Kuhn-Munkres](#) linear programming algorithm (a.k.a. the "Hungarian algorithm"). The input matrix to the Hungarian algorithm would be:

	Division 1	...	Division J
Worker 1	$v_{1,1}p_{1,1}$...	$v_{1,J}p_{1,J}$
...
Worker I	$v_{I,1}p_{I,1}$...	$v_{I,J}p_{I,J}$

A.2 Description of Deferred Acceptance Algorithm

For readers unfamiliar with the deferred acceptance algorithm, this section provides an overview. This algorithm, also known as Gale-Shapley algorithm ([Gale and Shapley, 1962](#)), is used in various matching markets. For management scholars, understanding this algorithm is beneficial as it opens up avenues to optimal decision-making, resource allocation, and efficient organizational structure.

The deferred acceptance algorithm aims to solve a problem of matching where each side has preferences over potential matches. To simplify this, consider an internal assignment scenario where you have workers and divisions. Here's a step-by-step guide to the algorithm:

- 1) *Initialization*: Each worker proposes to their top-choice division.
- 2) *Rejections & Tentative Acceptances*: Each division tentatively accepts the one worker they like best, and rejects all other applicants.
- 3) *New Applications*: Workers that were rejected in the previous round apply to their second-choice division.

- 4) *More Rejections & Tentative Acceptances*: Each company now considers this round's applicants along with their current tentatively accepted applicant. They keep the one they like best (which may be the one from the previous round) and reject the others.
- 5) *Repetition*: Steps 3 and 4 are repeated until each job seeker is either tentatively matched with a company or has exhausted all options and remained unmatched.
- 6) *Finalization*: All tentative matches become final.

To speed the execution of the algorithm, most programs ask participants to submit a full set of rankings on both sides so that the steps can proceed automatically based on the rankings.

The deferred acceptance algorithm has some notable properties:

- **Stability**: The matching is "stable" in that no unmatched division:worker pair would both rather be matched with each other over their current match. This prevents a situation where a division and a worker want to discard their assigned match to pair with each other. This does not mean that every participant gets their top choice. It does mean that if someone does not get their top choice, it is because the top choice is paired with someone who the top choice ranked higher.
- **Optimality**: For the side that makes the first move (in this case, the job seekers), this algorithm yields the best possible average outcome from among all stable matchings.
- **Deterministic**: For a given set of preferences, the algorithm always produces the same result.

For additional reading, [Roth \(2008\)](#) contains a discussion of the history, theory and practice around the deferred acceptance algorithm.

A.3 Proofs

A.3.1 Proof for Lemma 1

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 ($\bar{R} = (1 - G(\underline{u}))^2$).*

Proof. We begin with the Definition 5 for the definition of $\mathbb{E}[\pi]$, which says that $\mathbb{E}[\pi(a)] = \sum_{i=1}^I \sum_{j=1}^J \alpha_{ij}(a) v_{ij} P(i \leftrightarrow j \text{ acceptable})$.

Note that because of our I.I.D. assumption, $P(i \leftrightarrow j \text{ acceptable}) = (1 - G(\underline{u}))^2$. We call $\bar{R} = (1 - G(\underline{u}))^2$, note that \bar{R} is a constant.

Because \bar{R} is a constant, we can factor it out of the summation. The firm-dicted allocation a that maximizes expected payoffs maximizes:

$$\mathbb{E}[\pi(a)] = \bar{R} \underbrace{\sum_{i=1}^I \sum_{j=1}^J \alpha_{ij}(a) v_{ij}}_{V(a), \text{ Definition 1}} \quad (7)$$

We know that $V(a)$ (Definition 1) is maximized by the [Kuhn-Munkres](#) algorithm. As such, $\mathbb{E}[\pi(a)]$ is maximized by $\bar{R}V(a^*)$ where a^* is the allocation that maximizes $V(a)$ (Definition 1) using [Kuhn-Munkres](#). \square

A.3.2 Proof for Lemma 2

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.

Proof. We use the fact that under our I.I.D. assumption, DA selects all allocations with equal probability. Let this probability equal $\bar{\alpha}$. In the set of all allocations A (size $|A| = N!$) there are $(N-1)!$ allocations in which i is matched to j . As such, $\bar{\alpha} = (N-1)!/N!$ and so $\bar{\alpha} = 1/N$, which is a constant.

In addition, the ex-ante (unconditional) $P_{DA}(i \leftrightarrow j \text{ acceptable})$ for deferred acceptance is also constant (since the CEO does not know preferences ex-ante). We can then apply these facts to the definition of $\mathbb{E}[\pi]$ (Definition 5), which says that

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

We can factor both constants — both $\bar{\alpha}$ and $P_{DA}(i \leftrightarrow j \text{ acceptable})$ — out of the summation. As such, the expected payoff for DA is:

$$\begin{aligned} \mathbb{E}[\pi_{DA}] &= \bar{\alpha} \sum_{i=1}^I \sum_{j=1}^J v_{ij} P_{DA}(i \leftrightarrow j \text{ acceptable}) \\ &= \frac{1}{N} \sum_{i=1}^I \sum_{j=1}^J v_{ij} P_{DA}(i \leftrightarrow j \text{ acceptable}) \\ &= \underbrace{P_{DA}(i \leftrightarrow j \text{ acceptable})}_{R_{DA}, \text{ Expected DA Retention}} \cdot \underbrace{\frac{1}{N} \sum_{i=1}^I \sum_{j=1}^J v_{ij}}_{\bar{V}, \text{ Mean of all } V(a)\text{'s}} \end{aligned} \quad (8)$$

We finally show that the expression labeled \bar{V} in the equation above is equal to the average $V(a)$ over all possible a s. That is:

$$\frac{1}{N} \sum_{i=1}^I \sum_{j=1}^J v_{ij} = \frac{1}{N!} \sum_{a \in A} V(a) \quad (9)$$

Starting with the LHS, note that in the set of all allocations A , there are $(N-1)!$ different ways to match a given i and j .

$$\frac{1}{N!} \sum_{a \in A} V(a) = \frac{1}{N!} (N-1)! \sum_{i=1}^I \sum_{j=1}^J v_{ij} = \frac{1}{N} \sum_{i=1}^I \sum_{j=1}^J v_{ij} \quad (10)$$

□

A.3.3 Proof for Lemma 3

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

Proof. The proof of Proposition 2 below implies that the distribution of utilities over DA matches has a lower hazard rate than the unconditional distribution of utility. This implies that the distribution of utilities over DA matches first-order stochastically dominates the unconditional distribution and so has a higher retention rate for all \underline{u} such that $G(\underline{u}) < 1$. □

A.3.4 Proof for Proposition 1

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

Proof. If workers are unspecialized under Definition 6, then v_{ij} equals a constant v_i for the worker i , and $V(a) = \sum_{i=1}^N \sum_{j=1}^N \alpha_{ij}(a) v_{ij} = \sum_{i=1}^N v_i = V$ for all a . The most productive match (V_{FD}) is equal to the average match (\bar{V}). All the mass in \mathcal{V} is concentrated in a single point, and the LHS of Equation 6 is equal to 1. Lemma 3 shows that the RHS of Equation 6 is equal to or greater than one. □

A.3.5 Proof for Corollary 1

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

Proof. For firm-dictated assignment to outperform delegation, the LHS of Equation 6 needs to exceed the RHS. The right-hand side of Equation 6 will be greater than or equal to 1 because of Lemma 3. If workers are not specialized, the LHS of Equation 6 will equal 1 under Proposition 1. \square

A.3.6 Proof for Proposition 2

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

The unconditional quit probability on either side of the market is $G(\underline{u})$. This is increasing in \underline{u} . Our proof below shows that the returns to DA are increasing in \underline{u} . Then, all increases in $G(\underline{u})$ can be shown to be equivalent to an increase in \underline{u} .

Suppose that the utilities from DA matches are distributed according to $H(u)$ for the proposing side and $L(u)$ on the receiving side with densities $h(u)$ and $l(u)$, respectively. Then the retention rate from matching with DA, R_{DA} , is given by: $R_{DA}(\underline{u}) = (1 - H(\underline{u}))(1 - L(\underline{u}))$. The unconditional expected retention rate is $\bar{R}(\underline{u}) = (1 - G(\underline{u}))^2$.

Differentiating with respect to \underline{u} implies that the ratio $\frac{R_{DA}(\underline{u})}{\bar{R}(\underline{u})}$ is increasing if:

$$\frac{g(\underline{u})}{1 - G(\underline{u})} \geq \frac{1}{2} \left(\frac{h(\underline{u})}{1 - H(\underline{u})} + \frac{l(\underline{u})}{1 - L(\underline{u})} \right).$$

This says the relative retention benefits of DA are increasing if the unconditional hazard rate is greater than average hazard rate on the proposing and receiving sides. This is a form of stochastic ordering called the hazard rate ordering which is stronger than stochastic dominance.

The proof will require the following result.

Lemma A1. *For any absolutely continuous distribution $F(\cdot)$ with density $f(\cdot)$ such that $F(u) < 1$:*

$$\frac{f(u)}{1 - F(u)} \geq \frac{\frac{N!}{(r-1)!(N-r)!} F(u)^{r-1} (1 - F(u))^{N-r} f(u)}{1 - F_r(u)}$$

where $F_r(u)$ is the CDF of the r^{th} out of N order statistic for $r > \frac{N}{2} + 2$ and the right-hand side of the inequality is the hazard rate of $F_r(u)$.

Proof. We will use a proof by contradiction. Suppose the lemma is false, so:

$$\frac{f(u)}{1 - F(u)} < \frac{\frac{N!}{(r-1)!(N-r)!} F(u)^{r-1} (1 - F(u))^{N-r} f(u)}{1 - F_r(u)}.$$

Cross-multiplying and canceling the shared terms implies:

$$\frac{1 - F_r(u)}{1 - F(u)} < \frac{N!}{(r-1)!(N-r)!} F(u)^{r-1} (1 - F(u))^{N-r},$$

$$1 < F_r(u) + \frac{N!}{(r-1)!(N-r)!} F(u)^{r-1} (1 - F(u))^{N-r+1}.$$

Now, note that:

$$F_{r-1}(u) = F_r(u) + \frac{N!}{(r-2)!(N-r+1)!} F(u)^{r-1} (1 - F(u))^{N-r+1}.$$

This shows that $F_{r-1}(u)$ must be greater than the previous expression because they are equal except the second term is multiplied by $(r-1)(N-r+1) > 1$. Therefore:

$$1 < F_r(u) + \frac{N!}{(r-1)!(N-r)!} F(u)^{r-1} (1 - F(u))^{N-r+1} \leq F_{r-1}(u).$$

Because $r > \frac{N}{2} + 2$, $F_{r-1}(u) < F(u)$ because the distribution of an above median order statistic first order stochastically dominates the unconditional distribution. The above inequalities imply $F(u) \geq 1$ but we assumed $F(u) < 1$ so this is a contradiction. \square

Lemma A1 shows that the hazard rate of an order statistic above the median is lower than the unconditional hazard rate. We will show that both $H(u)$ and $L(u)$ are both distributions for above median order statistics.

Let Q_i be worker i 's rank of her assigned division in a worker-proposing stable matching, \mathcal{M} . Define $\mathcal{Q} = \max_i Q_i(\mathcal{M})$. \mathcal{Q} is the ranking of the least preferred match of any worker in this matching. Pittel (1992), Theorem 6.1 shows that $P(\mathcal{Q} \leq (2+a)\log^2 N) \geq 1 - O(N^{-c})$ for all $c < c(a)$ where $c(a) = 2a[3 + (4a+9)^{1/2}]^{-1}$.¹ In words, this says it is almost certain that the least preferred match of a worker in an $N \times N$ random matching market using working-proposing deferred acceptance will be ranked less than $(2+a)\log^2 N$. This suggests the worker who receives the least preferred match in the entire market is almost certain to match with utility draw at least in the p_N percentile, where $p_N = \frac{N - (2+a)\log^2 N}{N}$. For a fixed a , this is greater than $\frac{1}{2}$ for sufficiently large N . For example, if $a = 0$ this is greater than $\frac{1}{2}$ if N is about 75. If $a = 1$ this is greater than $\frac{1}{2}$ if N is about 150.

This result implies a lower bound on $H(u)$ is the distribution of an above median order statistic in sufficiently large markets.

All that is left to the prove the result is to show that:

$$\frac{g(u)}{1 - G(u)} \geq \frac{l(u)}{1 - L(u)}.$$

¹Pittel (1992) assumes all matches are acceptable. In our setting, the sub-market with only acceptable matches is expected to be of size $(1 - G(u))N$. Therefore, this restriction just re-scales the market size by the constant $(1 - G(u))$

While the receiving side of the market fairs worse than the proposing side, they do better than random because they can choose their more preferred offer when they receive multiple offers. Pittel (1989), Theorem 2 shows that the number of proposals made by the proposing side of the market converges in probability to $N \log(N)$. Therefore, an average member of the receiving side of the market is expected to receive $\log(N)$ proposals. Each of these proposals allows for another I.I.D. preference draw. As a result, $L(u)$ is distributed like an above median order statistic. The result therefore follows from Lemma A1.

A.3.7 Proofs for Corollaries 2, 3 and 4

Corollaries 2, 3 and 4 are applications of Proposition 2.

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 .*

Proof. We need to show that for any $u \in \mathbb{R}$ and $\mu - K \sim G'$, there exists another constant ϵ such that $G'(u) = G(u + \epsilon)$. Then the statement would be true under Prop 2 and Corollary 3.

If G first order stochastic dominates G' , then $G'(u) \geq G(u)$ for all u . Because G is a continuous CDF between zero and one and 1 and is increasing in its argument, you can add a positive constant ϵ to u until $G'(u) = G(u + \epsilon)$. As such, this statement is true under Prop 2 and Corollary 3. \square

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

Note: This was shown as part of the proof for Proposition 2.

Proof. If \bar{u} improves, then $G(\bar{u})$, the probability of quitting has increased because G is increasing in its argument. DA is more attractive by Proposition 2. \square

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 .*

Proof. We show that for any $u \in \mathbb{R}$ and G' (a mean preserving spread of the absolutely continuous distribution G), there exists a $\epsilon \in \mathbb{R}$ such that $G'(u) = G(u + \epsilon)$. Then the statement would be true under Proposition 2 and Corollary 3.

If G' is a mean preserving spread of G , then the CDFs of G and G' intersect at a single point (Diamond and Stiglitz, 1974). Call this single crossing point $\hat{\mu}$ such that $G'(\hat{\mu}) = G(\hat{\mu})$. Below this point, $G'(u) \geq G(u)$, and above this point $G(u) \geq G'(u)$ (Diamond and Stiglitz, 1974).

If the outside option \underline{u} is not too high ($\underline{u} < \hat{\mu}$), then $G'(\underline{u}) \geq G(\underline{u})$. Because both G and G' are continuous and lie within $[0,1]$, and because G is increasing in its argument, a positive constant ϵ can be added to \underline{u} until $G'(\underline{u}) = G(\underline{u} + \epsilon)$. As such, this statement is true under Prop 2 and Corollary 3. \square

A.3.8 Proof for Proposition 3

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.*

Proof. Because workers and divisions have strictly broadly aligned preferences, $\mu_j^i = u_i^j = V(a)$ if matching together will yield overall match a given all other workers' and divisions' preference reports. V_{FD} is the output of the optimal match. Therefore, it must be the case that $V_{FD} \geq V(a)$ for all $a \in 1, \dots, N!$. Suppose all workers except worker i rank their optimal match first. Then worker i 's best response is to also rank her optimal match, say j^* first because $\mu_j^i = V_{FD} \geq V(a)$ for all $a \in 1, \dots, N!$. The same logic applies to division j . This proves that the optimal assignment is a Nash equilibrium.

Now, suppose there is another Nash equilibrium that yields output $\tilde{V} < V_{FD}$. If all workers except i rank their match corresponding to this alternative equilibrium first, then it is a best response for i to also rank this alternative equilibrium first by our assumption that it is a Nash equilibrium. The same logic applies to divisions. Together, these results prove that there are possibly Nash equilibria that yield the same output as firm-dictated assignment but there are possibly others that do not. \square

A.3.9 Proof for Proposition 4

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

Proof. This follows from the examples in Appendices A.6 and A.7. \square

A.3.10 Proof for Corollary 5

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.*

Proof. Narrow alignment implies deferred acceptance will yield positive assortative matching. Assortative matching maximizes output when production is super-modular (Becker, 1973). \square

A.4 Noisy CEO Beliefs

Assumption 6 (Noisy CEO Beliefs). *Suppose that match-specific productivities v_{ij} are drawn I.I.D. from a distribution F with mean \bar{v} . CEO's do not directly observe each v_{ij} draw from F , but instead observe each v_{ij} with noise; i.e. they observe $\tilde{v}_{ij} = v_{ij} + \epsilon_{ij}$, where $\epsilon_{ij} \sim \mathcal{N}(0, \sigma^2)$ and drawn independently for each possible match.*

Proposition A1. *As the noise in CEO's observations of match productivity (σ^2) increases, the relative productivity gain from firm-dictated assignment decreases.*

Proof. When implementing matches with firm-dictated assignment, the CEO will select the match with the highest expected productivity because quit probabilities are identical across matches. Let $\tilde{v}_{ij} = v_{ij} + \epsilon_{ij}$ where ϵ_{ij} are I.I.D. $\mathcal{N}(0, \sigma^2)$.

The CEO perceives that an arbitrary match a has total productivity:

$$\tilde{V}(a) = V(a) + \sum_{i=1}^N \sum_{j=1}^N \alpha_{ij}(a) \epsilon_{ij},$$

where $\alpha_{ij}(a)$ is equal to 1 if worker i is matched to division j in match a . Because ϵ_{ij} are I.I.D. across individual worker-division matches, the sum of the ϵ s is distributed $\mathcal{N}(0, N\sigma^2)$.

Suppose $V(a) > V(b)$. The CEO correctly perceives the relative ordering of these two matches if:

$$\begin{aligned} V(a) + \sum_{i=1}^N \sum_{j=1}^N \alpha_{ij}(a) \epsilon_{ij} &> V(b) + \sum_{i=1}^N \sum_{j=1}^N \alpha_{ij}(b) \epsilon_{ij}, \\ V(a) - V(b) &> \sum_{i=1}^N \sum_{j=1}^N (\alpha_{ij}(b) - \alpha_{ij}(a)) \epsilon_{ij}. \end{aligned}$$

Notice that worker-division matches that are included in both match a and b do not contribute to the perceived difference in match productivity. Suppose there are K shared matches and $N-K$ distinct matches. Then $\sum_{i=1}^N \sum_{j=1}^N (\alpha_{ij}(b) - \alpha_{ij}(a)) \epsilon_{ij} \sim \mathcal{N}(0, (N-K)\sigma^2)$.

With this, we can calculate the probability that the CEO correctly perceives $V(a) > V(b)$ as $\Phi\left(\frac{V_a - V_b}{\sqrt{2(N-K)\sigma^2}}\right)$ which is decreasing in σ . Therefore, the CEO is less likely to correctly rank any pair of matches when productivity is observed with more noise. With firm-dictated assignment, the CEO will select the match with the highest perceived productivity. A necessary condition for selecting the best match is correctly perceiving that it yields higher productivity than every other match. We have shown that this probability is decreasing in σ . The same logic can be applied to the second-best match and so on. The benefit of FD is decreasing in σ . At the extreme, the best match will converge to the average match as noise diverges.

□

A.5 Replacement and Quitting Costs

Suppose every successful match has a benefit of $\frac{c}{N}$ in addition to its productivity v_{ij} . This additional benefit $\frac{c}{N}$ can be interpreted as a quitting cost (i.e., it is lost if the match is unsuccessful).

Lemma A2 (Quitting Costs). *Higher quitting costs for the firm attenuate the relative benefits of FD over DA, $\frac{V_{FD}}{V_{DA}}$.*

This result shows DA is relatively more appealing in tight labor markets where quitting costs may be high for the firm.

Proof. To see this, assume every successful match has a benefit of $\frac{c}{N}$ in addition to its productivity v_{ij} . We can interpret $\frac{c}{N}$ as a quitting cost because it is lost if the match is unsuccessful. The ratio $\frac{V_{FD}}{\bar{V}}$ is decreasing in c :

$$\frac{\partial}{\partial c} \left[\frac{V_{FD} + c}{\bar{V} + c} \right] = \frac{\bar{V} - V_{FD}}{(\bar{V} + c)^2}. \quad (11)$$

This is negative whenever $V_{FD} > \bar{V}$. Therefore, increasing the costs of quits makes DA relatively more appealing than FD. □

A.6 Narrow Alignment not Sufficient

We show this with a counterexample. Consider a 2×2 firm where the CEO's v_{ij} matrix is:

	Division A	Division B
Worker 1	x	1
Worker 2	1	0

where $1 < x < 2$. The optimal assignment is $\{(\text{Worker 1, Division B}), (\text{Worker 2, Division A})\}$, which yields two for the CEO.

Suppose that workers and divisions rank their match partners based on v_{ij} (narrow alignment). Then Worker 1 prefers Division A, who reciprocates (this is the consequence of narrow alignment). In a preference-respecting setup, the match will be $\{(\text{Worker 1, Division A}), (\text{Worker 2, Division B})\}$, which yields an output of x (less than the optimal output of two). The expected output of a delegated mechanism is lower, despite workers and division being narrowly aligned.

A.7 “Misaligned” Participants Leading to Optimal Assignments

Consider a 2×2 firm where the CEO’s v_{ij} matrix is given by:

	Division A	Division B
Worker 1	1	4
Worker 2	$\frac{1}{2}$	2

Assume preferences are *misaligned*, so workers and divisions both prefer to be in a match where they are individually least productive but find all matches acceptable. The CEO’s optimal assignment is $\{(\text{Worker 1, Division B}), (\text{Worker 2, Division A})\}$, which yields a total output of 4.5. If the organization delegates the decision to the market participants using a mechanism which respects agents’ preferences such as DA, it will achieve this assignment despite the nominally misaligned preferences of the participants.

A.8 Correlated Preferences: Same-Side and Cross-Side

Preferences could also depart from I.I.D. by exhibiting correlations among the workforce.

Vertical Preferences. In many real world settings, agents on one side of the market agree on the ranking of agents on the other side (vertical preferences). Managers may all agree on the best workers, and workers may agree on the best managers/divisions. Vertical preference assumptions appear in many academic papers.² Vertical preferences are also implied by the special case considered in Corollary 5 where the workforce is strictly narrowly aligned, and production technology is supermodular.

To introduce vertical preferences in our setting, we assume that the CEO understands which sides have vertical preferences (one or both), and that any remaining sides have I.I.D. preferences. However, the CEO does not know exactly which options are ranked 1st, 2nd, 3rd, (... etc.) by the vertical side(s).³ We begin by assuming only one side of the market has vertical preferences.

Proposition A2 (Vertical Preferences). *If one side of the market has vertical preferences, DA does not increase that side’s retention.*

Proof. If one side has vertical preferences then regardless of which of the $N!$ matches is selected, exactly one agent on that side will match with her 1st, 2nd, ..., N th most preferred division because this is a one-to-one match.

²For example, [Agarwal and Diamond \(2017\)](#) consider the econometric identification of preferences in a two-sided market under vertical preference assumptions.

³In practice, this may be easy for a CEO to find out. If the CEO learned the vertical preferences, this would decrease the value of DA, insofar as it would allow the CEO to incorporate information about quitting without using DA.

Without loss of generality, suppose the worker side has vertical preferences. Let $P_{a,(i)}$ denote the expected retention of the worker who matches with her i^{th} ranked choice in match a . The overall retention rate of workers will be $\bar{P}_a = \frac{1}{N} P_{a,(i)}$.

But the average across all order statistics is just the unconditional average so $E[\bar{P}_a] = G(\underline{u})$ is the unconditional retention rate. In expectation, there is no variation in retention among the $N!$ possible assignments a with regards to the vertical side's quit rates. \square

The intuition of this result is that if preferences on one side of the market are vertical, then all participants on that side have the same first choice, second choice, etc. No matter how these agents are matched or who proposes, exactly one agent will match with her k^{th} choice for all k , so FD and DA have identical expected retention rates. If the other side's preferences are I.I.D., there may still be benefits from using DA from retaining that side, particularly if the I.I.D. side proposes. These benefits may be high enough to outweigh the coordination benefits of firm-dictated assignment. However, the DA retention benefit will be lower (versus if both sides had I.I.D. preferences).

Semi-Vertical Preferences. Proposition A2 shows that perfectly vertical preferences eliminate the potential benefits of DA. However, most preferences are not perfectly vertical; they may be highly correlated but not completely. Are the retention benefits of DA monotonically decreasing in workers' preference correlation? Although we do not have analytic results on this question, we present simulation results in Figure C3. The figure suggests that the DA retention rate increases initially as preferences become slightly correlated, but converges to the unconditional retention rate as preferences become highly correlated.⁴

Cross-Sided Correlations. Preferences can also be correlated across sides. Workers' and managers' preferences may exhibit a high degree of reciprocity. If preferences are correlated across sides, but are otherwise I.I.D. (i.e., no vertical preferences), then the results in our I.I.D. model all hold. None of our results from this model rely on uncorrelatedness across sides. They are mainly driven by the lack of correlation between workforce preferences and the production technology.⁵

⁴As we show in Appendix Figure C3, the initial increase in overall retention rates (on the left-hand side of Figure C3) is a result of the receiving side of the market receiving a better match. But as preferences become more correlated, overall retention declines because more workers are matched with less preferred divisions. Appendix Figure C3 contains the overall retention rate, as well as retention rates broken out by the workers and divisions side.

⁵The addition of cross-sided correlations does not change expected retention rate under DA, and thus Equation 6 is the same. This is partly because of our formulation of the problem, in which the CEO's penalties are the same irrespective of if an agent quits, or if the agent's match partner quits, or both quit. If the penalty for both quitting is higher, then cross-sided preferences would result in a higher degree of both-sides quitting than in our standard I.I.D. setup, and the benefits of DA would be higher (since the CEO gets extra credit for avoiding both sides quitting).

A.9 Equivalence between Random Utility and Random Outside Options

In this appendix, we demonstrate that the random utility model we use in our main analysis (I.I.D. random match utility with common outside option) generates the same key results compared to an alternative model featuring both I.I.D. random outside options as well as I.I.D. random outside options.

In other words, adding I.I.D. random outside options does not change the results in any of the propositions, lemmas, or corollaries of the private-value model in Section 3.2-3.5. These results are mostly directional, and the addition of I.I.D. does not change the directions. Below, we show the isomorphism between these two models.

In our main analysis, we assume that workers' utilities, μ_j^i , and divisions' utilities, u_i^j , are independently and identically distributed according to an absolutely continuous distribution $G(\cdot)$. Both workers and divisions have a common outside option, \underline{u} . Under these assumptions, a match with division j is acceptable to worker i if $\mu_j^i \geq \underline{u}$. Similarly, a match with worker i is acceptable to division j if $u_i^j \geq \underline{u}$. The probability that either a worker or division finds a match acceptable is $1 - G(\underline{u})$.

Now, suppose that we add I.I.D. outside options. Outside options, \underline{u} , are independently and identically distributed according to $W(\cdot)$. A match remains acceptable to an agent if that agent's utility from the match exceeds her outside option. The probability that a worker finds a match acceptable is therefore:

$$\begin{aligned} P(\mu \geq \underline{u}), \\ P(\mu - \underline{u} \geq 0). \end{aligned}$$

The difference $\mu - \underline{u}$ is the convolution of μ and \underline{u} which itself follows a distribution G' . The probability that a worker finds a match acceptable is now $1 - G'(0)$. This is simply an instance of our baseline model where the probability is $1 - G(\underline{u})$ (simply let $G = G'$ and $\underline{u} = 0$).

Compared to the model without random outside options, adding the random outside options (holding the distribution of agent utilities constant) would change the quit probabilities to something numerically different. However, the propositions, lemmas and corollaries in our main results were not based on specific numerical realizations. These results were directional, and were based on the structure of the CEO's optimization problem and constraints. The results still apply with the addition of I.I.D. outside options because the new model can be expressed as an instance of the main one.

A.10 Correlation between Worker Productivity and Outside Options

Section A.9 demonstrates that the random utility model with fixed outside options that we use in the main paper is equivalent to an alternative model with random utilities and outside options. In this section, we present a different type of model featuring heterogeneous outside options. Rather than presenting I.I.D. outside options, we allow correlation between a worker’s productivity and outside options.

To develop this type of model, we must adapt our setup to one where workers are vertically differentiated on productivity. Suppose that workers’ productivities are not match-specific, and instead each worker i has a fixed productivity v_i that is constant across all possible assignments. Let \bar{v} be the average productivity across all workers. For two workers i and j , if $v_i > v_j$ means that i ’s productivity dominates j ’s across all assignments.

Based on this setup alone, we know that Deferred Acceptance will dominate Firm-Dictating because there is no match-specific productivity. The workforce is completely “unspecialized” (Definition 6) and thus DA will dominate (Proposition 1). Nevertheless, we are interested in how the performance of DA changes as outside options become more (or less) correlated with the productivity v_i term.

To show this, we impose the assumption that at least one match is acceptable to every worker. Without this assumption, it is possible that outside options become so good that the most productive workers quit no matter what; in this case, they will leave under both DA and FD.

To show the effect of outside options that are correlated with productivity, we allow worker’s utilities over matches, μ_j^i , to be I.I.D. (as in our main model). Suppose that worker productivity, and worker i ’s outside option, u_i , are jointly normally distributed:

$$\begin{pmatrix} v_i \\ u_i \end{pmatrix} \sim \mathcal{N} \left(\begin{pmatrix} \bar{v} \\ 0 \end{pmatrix}, \begin{pmatrix} \sigma^2 & \rho \\ \rho & \sigma^2 \end{pmatrix} \right).$$

Divisions are narrowly aligned, and have vertical preferences over workers such that $u_i^j = u_i = f(v_i)$ for all j (where f is an increasing function). Therefore, all divisions have the same rank-order preferences over workers and desire workers in the order of their (vertically differentiated) productivity.

Simulation. This model is difficult to estimate analytically. Instead, we simulate the following model 10,000 times each for $\rho = -1, -0.9, \dots, 1$:

- $N = 100$,
- $V = \sum_i \sum_j \alpha_{ij} v_i \mathbb{1}(\mu_i^j \geq u_i)$,
- $\begin{pmatrix} v_i \\ u_i \end{pmatrix} \sim \mathcal{N} \left(\begin{pmatrix} 5 \\ 0 \end{pmatrix}, \begin{pmatrix} 1 & \rho \\ \rho & 1 \end{pmatrix} \right)$,

- $\mu_i^j \sim \mathcal{N}(0,1)$.

We have set the worker productivity distribution to have a mean of 5 and standard deviation of 1 so the firm will almost always prefer to retain a worker. As mentioned above, assume that all workers have at least one acceptable match. To this end, we calculate the difference between every worker's utility over their most preferred match, u_i^* and their outside option, \underline{u}_i . If this is negative for any workers, then those workers do not find any matches acceptable. To impose the assumption, we subtract the most negative $u_i^* - \underline{u}_i$ value (or zero if there are no negative values) from all utility draws.

Deferred Acceptance. Because divisions have vertical preferences over workers, deferred acceptance reduces to a serial dictatorship with the most preferred worker choosing first, the second most preferred worker choosing second, and so on.

Firm-Dictated Assignment. The firm knows each worker's productivity and would like to retain the most productive workers. The firm can anticipate who has relatively good or bad outside options because it observes productivities (v_i) and knows ρ . Unfortunately, the firm has no way of leveraging this information when dictating matches because it does not know workers' preferences. Therefore, the best the firm can do is randomly assign matches.


Results. The left panel of Figure C4 shows how overall productivity varies with ρ . Notably, DA dominates FD across all ρ values. This is because DA leverages workers' private information about their preferences so $\frac{R_{DA}}{R} > 1$ but there are no productivity differences across matches for FD to leverage so $\frac{V_{FD}}{V} = 1$. However, the simulation shows that the benefits of DA are increasing in the correlation between productivity and outside options, especially as the correlation becomes large.

The right panel of Figure C4 sheds light on what is driving this result. The figure shows how the retention rate of the top 10 workers varies with ρ when DA or FD is used. Regardless of whether DA or FD is used, the retention of the best workers declines with ρ . This is because outside options are correlated with productivity so more productive workers will have better outside options when ρ is more positive. When DA is used and the correlation is more positive, the best workers get their most preferred matches, so the retention rate declines more slowly. When FD is used, the firm has no way to reward more productive workers with better matches, so retention declines more quickly with ρ .

B Institutional Details & Data Appendix



B.1 Screenshots of Internal Marketplace (Mockups)

Figure B1: Profile Screen















































Notes. This figure portrays a mock profile screen of a manager.

Figure B2: Ranking Screen


EN ▼  Amy Smith ▼

Instructions
My Profile
My Positions
All Positions
Candidates

← **Role: Software Engineer (Hooli Banking)**

RANKED	UNRANKED	EXCLUDED
(1)  Lila 	 Jessica 	 Kenneth 
(2)  Derrick 	 Michael 	(X)  Jason 
(3)  Jennifer 	 John 	(X)  Gloria 
(4)  Anita 	 Ashley 	 Maria 
(5)  Isaac 	 Ryan 	 William 
(6)  Shawn 	 Cynthia 	 Victoria 
	 Abigail 	 Thomas 
	 Richard 	 Sarah 

[Home](#) | [Contact](#) | [FAQs & Help](#)

Notes. This figure portrays a mock ranking screen for a manager ranking potential workers in the deferred acceptance algorithm. The deferred acceptance algorithm uses agents’ ranked order preferences. The column on the left (“Ranked”) accepts inputs about the agent’s ranking over match partners. #1 is the top (preferred) rank, #2 is second, and so on. The webpage expanded so that the agent could rank as many options as they wanted.

The deferred algorithm can accommodate preferences over being unmatched ([Roth and Sotomayor, 1990](#); [Roth, 2008](#)). Agents may prefer to be unmatched rather than to accept certain partners they dislike. The rightmost column (“excluded”) allows agents in the setting to identify which match partners are below the “unmatched” threshold.

Finally, the deferred algorithm can also accommodate ties or indifferences ([Roth and Sotomayor, 1990](#); [Roth, 2008](#)). In our setting, agents often felt that certain options were better than being unmatched, but were otherwise “tied for last place.” The user interface used the middle column of Figure B2 for these options (“unranked”). Workers were told to use this column s directoryfor options the they liked beneath the “ranked” options, but were still better than “unmatched” (excluded).

B.2 Information Sources for Market Participants

Inside the company, there were several sources of information that market participants could use to inform their choices to rank prospective match partners. Although we do not claim users have perfect or unlimited access to information, it appears that our setting was relatively favorable in terms of information access. We mention four particular features below.

- 1) Open internal access to key data. The company in question granted workers relatively open access to internal information. The code and key engineering design documents used in the company's products were visible to all engineers and product managers (including all workers in our study). In addition, all workers enjoyed access to a company directory showing the job assignments of all workers, and the identities of future teammates for each project. Note that prospective teammates were effectively fixed from a worker's perspective, because only one worker on each team would move at a time. This creates certainty about the identity of prospective teammates in each match, and the ability to research these teammates.
- 2) Open reach-out culture. The firm also welcomed workers and managers from different parts of the company to reach out to each other and initiate dialogue or a meeting. In our setting, workers could reach out to a prospective manager and arrange a meeting to discuss an opportunity on the internal market. Managers accepted these meetings in part because they typically welcomed the opportunity to grow their team by adding a new member.
- 3) Detailed matching profiles and directory. As described in the main text, the internal matching application contained a detailed profile of managers and workers, including their background, how long they have been employed at the firm, and the open position. Each agent could not only learn who was on the other side of the market, but also who was on their own side.
- 4) Explicit guidance from company executives. In other companies, a worker might not know what assignments the executives wanted him/her to pursue. However in our setting, the user interface of the match profile explicitly informed each worker and manager about the firm's preferred matching for that user (and to what degree they approved). This was communicated in the form of the color-coded icons described in Section 5 ("How the Firm Used the Scores"). Workers and middle-managers were told what these icons meant ahead of time, so that workers who cared about the firm's preferences could follow these suggestions.

Finally, the match system described in the paper (using deferred acceptance) gave workers an incentive to ask for additional information they did not have access to. Workers in the market who felt uninformed still had to submit rankings that would affect their assignments. As a result, each agent had an incentive to gather information in order to gain a match reflecting their priorities.

Our impression is that these policies made the firm relatively open with regards to access to information (and incentives to gather it). These policies were adopted as part of a philosophy of sharing information between senior managers and the broader company (in order to gain potentially useful input from workers, as your previous point discusses). A side-benefit of these policies is that workers could more easily learn about other teams' projects and form impressions of them.

B.3 Data

Our data consist of four tables about participant characteristics, preferences, pairwise match scores and potential assignments under various matching regimes.

Participant Characteristics. For both workers and managers, we have a set of individual covariates. Some of the manager characteristics are associated with the job he/she is hiring for. Note that some workers and managers are observed multiple times, as they finish one assignment and enter another one over our sample period.

Participant Preferences. For each worker and manager, we have a complete set of ordinal preferences over the opposite side of the market. For workers who appear in the market multiple times, we have multiple sets of preferences. Note: We describe our preferences as “complete,” although a large portion of possible matches were left tied for last place.⁶ Given that our setting has more managers than workers, some managers will not be matched to workers. Non-matches are coded as the least desirable choice for participants based on their inputs.⁷

Pairwise Match Quality Scores. For all possible pairs of workers and managers, we have several numeric scores characterizing the match. The most important in our analysis is the *firm match quality score*. This is the firm's estimate of match quality based on the firm's goals. The firm would prefer the match that maximizes the sum of total pairwise scores across all assignments.

Actual and Simulated Assignments. Finally, our data includes the assignments generated by the worker-proposed match used by the company. We have also re-run this algorithm 50 independent times, resolving indifferences and ties randomly, to generate a distribution of potential assignments. In addition, we also use the preference data to generate counterfactual assignments based on a variety of other assignment methods. Specifically, we examine the managers-propose DA algorithm which is nearly always identical to the workers-propose. We also examine completely random assignments, as well as random serial dictatorship ([Abdulkadiroğlu and Sönmez, 1998](#)) led by the workers and

⁶For example, if a worker ranked seven out of a total of ten managers, we code the three non-ranked managers as choice 8. We also run a version where we randomize participant preferences among unranked choices.

⁷As described above, all participants were given the option to rank “prefer to be unmatched” ahead of some least-preferred matches. No participant chose to utilize this option.

managers (we call these the “workers’ draft” and the “managers’ draft”). Our goal using these latter two is to help us evaluate non-parametrically how much our results are driven by workers and managers preferences by examining matching based only on one side’s preferences. We similarly run each of these algorithms 50 times, resolving ties randomly. When running matches using participants preferences, we utilize the same random draws to break ties across all different assignment algorithms.

Finally, we calculate the firm-dictated match, which maximizes the firm’s match quality score (as described above). To calculate this, we utilize the aforementioned [Kuhn-Munkres](#) algorithm (aka the “Hungarian algorithm,” [Kuhn 1955](#)). This algorithm identifies the set of assignments which maximize a numeric objective. We apply this algorithm to find the set of worker/manager assignments that maximize the total firm match quality score (described above). Because this sometimes involves resolving ties between pairs with the same match quality, we run this algorithm 50 times. We also run a version of the [Kuhn-Munkres](#) algorithm that minimizes the total firm objective to provide a lower bound of the firm’s objective function in our data (while still assigning all workers are matched).⁸

C Additional Empirical Analysis

C.1 Validating the firm objective score

In Table [C1](#), we examine realized matches only. We study whether good outcomes for the firm were forecast by the firm-objective score. 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.

Before each job was posted on the market, the posting was reviewed ex-ante by executives and managers to establish “required” and “preferred” skills for the role. These requirements were expressed in a structured way using ≈ 200 skills. “Required” skills were regarded as essential, and “preferred” were useful but not strictly necessary. As an outcome measure of match quality, we report what percentage of these skills were assessed as either “performing” (i.e., meets expectations) or “mastery” (exceeds) for the worker in their assigned matches. We use the data about the realized outcomes the firm adopted from the DA algorithm that were driven by workforce preferences.⁹

The average match in our data has an outcome of 66%, meaning that only 66% of the “required” skills were assessed as performing or mastery. This reflects how much workers

⁸One possible way to minimize the firm’s objective might be to leave them all unmatched; our implementation of the objective-minimizer forbade this.

⁹Ideally, a validation of the firm objective score would use randomized assignments or an instrument. Although we do not have random assignments, we do know the data generating process used to form these matches (deferred acceptance), including the preference inputs into this procedure. As our previous results indicate, these preferences were largely uncorrelated with the firm’s (see results in Section [7.1](#)).

and managers ignored the guidance of the firm objective scores, and have rankings that are broadly uncorrelated with the firm objective score. We find that a single standard deviation increase in the firm objective score corresponds to a 25 percentage point increase in the percentage of required skills assessed at performing or mastery level. This is an increase of 0.88 standard deviations. Similarly, a one standard deviation increase in the firm objective score corresponds to a 28 percentage point increase in the percentage of *preferred* skills at performing or mastery (0.72 standard deviations).

Table C1: Percent of Skills Assessed at Performing or Mastery Level

	Percent Required for Role	Percent Preferred for Role	Percent Required for Role, σ	Percent Preferred for Role, σ
Objective Score (v_{ij}), σ	.25*** (.024)	.28*** (.057)	.88*** (.087)	.72*** (.15)
Mean Dep. Var	.66	.57	0	0
Worker FEs	Y	Y	Y	Y
Quarter/Submarket FEs	Y	Y	Y	Y
R^2	.98	.93	.98	.93
Observations	428	428	428	428

Notes: The unit of observation in this regression is the worker \times manager pair. Only realized matches are included. The regressions predict whether good outcomes for the firm were predicted by the objective score. Columns 1 and 3 examine percent outcomes between zero and one, and Columns 2 and 4 examine normalized outcomes.

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

C.2 Measurement of Preference Correlations

Same-Side Preference Correlations. To measure the degree of correlated preferences, we create a dataset containing all possible pairs of workers, and calculate the Kendall (B) rank-correlation τ for each pair of workers' preferences over the same set of managers. We later repeat this analysis from the manager's side, examining how correlated pairs of managers' preferences are over the same set of workers. The correlations in these analyses could in theory vary from -1 to +1.

Table C2 shows our results. For both workers and managers, the average correlation is around zero. This means that one worker's preferences are uninformative about another's; workers have independent uncorrelated views of the same set of managers. Our results suggest workers and managers have a highly idiosyncratic, personalized, and distinctive set of preferences that do not come into conflict with other members of the same side. Prop A2 suggests that this is good for deferred acceptance, insofar as the workforce as a whole can obtain high rankings without competing.

Table C2: **Preference Correlations: Same-side & Cross-side**

	Mean	Std.Dev	Min	P25	Median	P75	Max
Same Side (All Ranks):							
Worker preferences	-0.01	0.09	-0.97	-0.02	-0.02	-0.02	1.00
Manager preferences	-0.01	0.15	-1.00	-0.03	-0.02	-0.02	1.00
Same Side (Top Rank):							
Worker preferences	-0.00	0.25	-0.99	-0.12	0.00	0.12	1.00
Manager preferences	0.01	0.30	-1.00	-0.13	0.01	0.15	1.00
Cross-side:							
All worker/manager pairs	0.21						
Mutually-ranked pairs	0.53						

Notes: This table displays summary statistics of pairwise correlations on the same side of the market, and across sides. Appendix C.2 contains additional details about how these numbers were calculated.

Cross-Sided Reciprocity. Our results above suggest that workers and managers’ tastes are both highly idiosyncratic and personal. We now turn to cross-sided preference correlations. If a worker (idiosyncratically) likes a manager, does the manager like the worker in return (and vice versa)? We measure this by calculating the Kendall (B) rank-order τ for cross-sided preferences.

Our results in the bottom of Table C2 suggest slightly stronger positive correlations (compared with our same-sided results), but still overall moderate correlations in cross-sided preferences. Among all preferences, the correlation is 0.20 (of a possible range of -1 to +1). Among pairs who ranked each other, the correlation is higher (0.53).

C.3 Misalignment or Coordination Failure?

As discussed in our theoretical framework, a broadly aligned workforce could settle on a bad equilibrium (among multiple possibilities). If DA produces low match quality, this type of coordination failure could be at fault. Alternatively, workers and managers may simply *not* be broadly aligned. To study these possibilities, we examine whether workers and managers make firm-optimal choices, conditional on what other players are doing.

We re-run the deferred acceptance algorithm over 450K times. Each time, we replace one agent’s submitted preferences with a randomly-drawn preference profile (holding all other preferences constant as they were submitted). We place our results in the Table C3 below.

Table C3: **Conditional Alignment**

	DV = Firm's Match Quality Score (Normalized)	
Submitted Preferences	.016* (.0086)	.013* (.0066)
Sample	Workers	Managers
R^2	.000015	9.3e-06
Observations	308,160	372,240

Notes: For each worker and manager in our sample, we replace their submitted preferences with a randomly drawn preference profile, one participant at a time. We then re-run the deferred acceptance algorithm for the entire market, using the actual preferences for all workers except the one with whose preference we replaced. We repeat this thirty times for each worker and manager. We then examine average match quality for the firm, for the randomly drawn preferences as well as the preferences actually submitted by the workers and managers. Because other participants' reports include some ties, we evaluate each preference profile across thirty random ways of randomly breaking ties.

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

C.4 Beliefs about Other Players' Actions

Table C3 suggests that agents in our setting are not making better-than-random choices for the firm's match quality, conditional on what other agents are doing. This raises the question: What set of beliefs about other players' actions would justify the submitted preferences as being good for the firm?

To study this question, we re-run the deferred acceptance algorithm multiple times iteratively. Each time, we focus on one focal worker's (or manager's) beliefs. For each focal worker, the simulation proceeded in four steps:

- 1) We draw a random set of preferences for all other players.
- 2) We measured how similar these random preferences were to the other players' actual preferences (i.e., the accuracy of these beliefs) by measuring rank correlation between the randomly-drawn beliefs and the actual rank-ordered lists. This produces a number between -1 and 1.
- 3) We draw 7 randomly selected alternative profile preferences for the focal worker.
- 4) Finally, we measure whether the worker's actual submitted preferences result in better (or worse) outcomes for the firm than the randomly chosen ones (holding the other agents preferences constant as in Step 1). We subtract the firm payoffs in the player's actual rank-ordered list from the average firm payoffs from randomly drawn preferences in Step 3).

We ran the four steps above ≈ 170 times per agent (breaking ties randomly when they appeared in actual preferences). This required running DA over 1M times under different permutations. We summarize the output in a dataset containing $\approx 20K$ observations,

where each observation is an agent's possible belief (randomly drawn), and a) how accurate each possible belief is, as well as b) whether each worker's options were helpful for the firm (had the beliefs been true), calculated as in Step 4).

We study the relationship between a) and b) in the table below. In Table C4 below, we can reject the hypothesis that there exist no beliefs that would justify the observed behavior as better-than-random ($p < 0.005$). About one third of the randomly-generated beliefs would justify at least one agent's ranked-lists as better-than-random, although most of these benefits are very small improvements over randomly chosen lists. In Table C5, we show that these payoffs are unrelated to the accuracy of these beliefs. Our coefficients do not allow us to reject zero correlation, and our standard errors are tight enough to show no economically large relationships.

Table C4: Are there Beliefs that rationalize the rankings for the firm?

	Better than Random	Improvement above Random (std)
Constant	.36*** (.0075)	-4.8e-10 (.013)
Sample	All	All
R^2	0	0
Observations	37,283	37,278

Notes: The regressions above analyze randomly drawn beliefs about other players, and measure whether a focal player has acted good for the firm, conditional on those beliefs. For more detail see Section C.4.

Table C5: Accuracy of Beliefs that rationalize the rankings

	Better than Random			Improvement above Random (std)		
Accuracy of Beliefs about Workers' Rankings	-.00079 (.0025)	-.0013 (.0038)	-.002 (.0029)	-.0026 (.005)	-.0073 (.009)	.0006 (.005)
Accuracy of Beliefs about Managers' Rankings	.0021 (.0024)	.00069 (.0037)	.0037 (.0028)	-.0015 (.0051)	-.0058 (.0091)	.0022 (.0053)
Sample	All	Workers	Managers	All	Workers	Managers
R^2	.000022	8.7e-06	.00011	9.4e-06	.000057	.00001
Observations	37,283	17,120	20,163	37,278	17,117	20,161

Notes: The regressions above analyze randomly drawn beliefs about other players, and measure whether a focal player has acted good for the firm, conditional on those beliefs. The table above evaluates the accuracy of these beliefs. For more detail see Section C.4.

C.5 Common Beliefs

The analysis in Appendix C.4 above finds little evidence that workers' rankings can be justified as better-than-random for the firm by using correct beliefs about what other

players are doing. However, it is possible that the workforce shares a common, inaccurate belief about what other players are doing. Although this belief may be inaccurate, it may be shared among many people in the market.

To evaluate this hypothesis, we extend the simulations from Appendix C.4. In the simulations in Appendix C.4, the set of beliefs were randomized separately for each worker. In the analysis that follows, we generate 21 sets of random beliefs. Each set of beliefs consists of 46K pairwise rankings between workers and managers (23K on each side). We hold these set of beliefs constant across workers. For each belief, we study whether the observed rankings are better than random for all workers, holding the belief constant. This allows us to study whether the observed rankings are good for the firm using “belief fixed effects.” Certain beliefs may rationalize behavior for all agents, even if these beliefs are not accurate.

When we test this hypothesis, we find little evidence of belief fixed effects. The joint test for the existence of belief fixed effects has a p -value of 0.30, and an R^2 of less than 0.01.

Table C6: **Common Beliefs**

	Submitted Rankings Improvement over Random Rankings (std)	
Belief FEs	Y	Y
Perturbed Agent FEs		Y
R^2	.0014	.17
Observations	13,268	13,268
p -value, joint test of Belief FEs=0	.73	.77

Notes: The regressions above analyze randomly drawn beliefs about other players, and measure whether a focal player has acted good for the firm, conditional on those beliefs. The table above includes “belief fixed effects” to show whether certain beliefs rationalize the rankings for all players. For more detail see Section C.5.

C.6 Complementarity and Submodularity of Match Value Scores

The optimal match for the firm depends partly on the degree of supermodularity in the match production function (Becker, 1973). To analyze complementarity empirically, we estimate the relationship between match quality and average match quality score of each worker and manager in the match.

Table C7 shows results. We find that quality is *increasing in the difference* between worker quality and manager quality. Column 1 shows that a one standard deviation increase in the interaction between the worker’s average match quality (across all matches) and the manager’s decreases the firm’s match quality score by ≈ 0.1 standard deviations ($p=0.001$).

Column 2 shows that the firm's match quality score increases as the squared quality difference between workers and managers increases. This suggests the firm-dictated match will exhibit a degree of negative assortative matching.

A key feature of the match quality scores provides intuition for our results: The executives' match quality scores featured a maximum and a minimum. We scale these to zero and one in our empirical analysis. Because of the upper bound, the performance of the whole can be "less than the sum of the parts;" i.e., performance has an upper bound that cannot be exceeded, no matter how great the parts are.

Table C7: **Complementarity: All worker-manager pairs**

	Firm Objective Score (Norm)	
Worker Quality \times Manager Quality	-0.099***	
	(0.0064)	
(Worker - Manager Quality) ²		0.050***
		(0.0032)
Worker Fixed Effects	Y	Y
Manager Fixed Effects	Y	Y
R^2	0.85	0.85
Observations	23361	23361

Notes: The regressions in this table use all possible worker-manager matches. Column 1 regresses the firm match quality score on the interaction between worker and manager quality, while column 2 regresses the firm match quality score on the absolute value of the difference between worker and manager quality. Worker and manager quality are normalized to have mean 0 and standard deviation one. Both columns include fixed effects for workers and managers.

C.7 Variations of Match Quality Table

Table C8: Firm Match Quality Score by Matching Algorithm: All Matches

	Mean	Min	P25	Median	P75	Max
Firm-Dictated	0.66 (0.04)	0.00 (0.00)	0.52 (0.07)	0.75 (0.06)	1.00 (0.05)	1.00 (0.00)
Managers-Propose DA	0.52 (0.05)	0.00 (0.00)	0.29 (0.08)	0.57 (0.07)	0.77 (0.11)	1.00 (0.00)
Workers-Propose DA	0.52 (0.05)	0.00 (0.00)	0.29 (0.08)	0.57 (0.07)	0.77 (0.11)	1.00 (0.00)
Worker Draft	0.52 (0.05)	0.00 (0.00)	0.30 (0.08)	0.56 (0.07)	0.77 (0.11)	1.00 (0.00)
Manager Draft	0.52 (0.06)	0.00 (0.00)	0.28 (0.07)	0.57 (0.07)	0.76 (0.12)	1.00 (0.00)
Random Assignment	0.49 (0.06)	0.00 (0.00)	0.22 (0.06)	0.52 (0.07)	0.71 (0.12)	1.00 (0.00)
Match Quality Minimizer	0.32 (0.08)	0.00 (0.00)	0.08 (0.04)	0.25 (0.06)	0.45 (0.18)	1.00 (0.13)

Notes: This table displays the average firm match quality score for assignments generated by various algorithms. Worker (Manager) Draft refers to assignments generated by random serial dictatorship [Abdulkadiroğlu and Sönmez \(1998\)](#) led by workers (managers), while Match Quality Minimizer refers to matches generated from the Kuhn-Munkres algorithm that minimize the total firm objective (while still ensuring all workers are matched). Standard errors of the mean were bootstrapped by the quarter of the market, using one thousand bootstrapped samples per matching strategy. Because there were more managers on the market than workers, some managers were unmatched. This table includes all managers, labeling unmatched managers a match quality score of zero. All matching algorithms left the same number of managers unmatched, but differed in their composition. A similar table including only complete manager/worker pairs is accessible in [Table 4](#).

Table C9: Effects on Firm Objectives

	Firm objective score					
Deferred Acceptance	-0.14*** (0.021)	-0.14*** (0.020)	-0.14*** (0.021)			
Random assignment	-0.16*** (0.024)	-0.16*** (0.024)	-0.16*** (0.024)	-0.16*** (0.024)	-0.16*** (0.024)	-0.16*** (0.024)
Managers-Propose DA				-0.14*** (0.021)	-0.14*** (0.020)	-0.14*** (0.021)
Workers-Propose DA				-0.14*** (0.021)	-0.14*** (0.021)	-0.14*** (0.021)
R^2	0.821	0.464	0.918	0.821	0.464	0.918
Observations	129250	129250	129250	129250	129250	129250
Fixed effect	Worker	Manager	Both	Worker	Manager	Both
P-values:						
DA = Random	0.017	0.017	0.017			
DA (MP) = Random				0.017	0.017	0.017
DA (WP) = Random				0.017	0.017	0.017

Notes: This table displays the results of a regression of the firm match quality score on indicators for the match being generated through deferred acceptance or random matching. The excluded category is the firm-dictated match. Each regression includes robust standard errors clustered at the simulation-submarket level. The bottom of the table displays p-values that test the difference in coefficients on deferred acceptance versus random assignment. The average firm match quality score from the firm-dictated matching is 0.62.

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

C.8 Variations of Worker/Manager Preferences Tables

Table C10: Manager Ranking of Assignment by Matching Algorithm

	Mean	Min	P25	Median	P75	Max
Manager Draft	8.07 (0.84)	1.00 (0.00)	1.00 (0.00)	1.00 (0.31)	3.00 (1.30)	106.00 (23.05)
Worker Draft	8.16 (0.77)	1.00 (0.00)	1.00 (0.00)	1.00 (0.05)	2.00 (1.29)	106.00 (16.61)
Managers-Propose DA	8.20 (0.84)	1.00 (0.00)	1.00 (0.00)	1.00 (0.00)	2.00 (1.35)	106.00 (16.47)
Workers-Propose DA	8.20 (0.85)	1.00 (0.00)	1.00 (0.00)	1.00 (0.00)	2.00 (1.34)	106.00 (16.47)
Random Assignment	8.98 (0.92)	1.00 (0.00)	2.00 (0.00)	3.00 (0.47)	5.18 (1.43)	76.68 (16.47)
Firm-Dictated	9.13 (0.93)	1.00 (0.00)	2.00 (0.00)	3.00 (0.51)	5.32 (1.53)	78.28 (22.96)
Match Quality Minimizer	9.64 (0.89)	1.00 (0.00)	2.00 (0.11)	3.00 (0.48)	6.00 (1.53)	104.08 (22.96)

Notes: This table displays how managers ranked their assignments. Worker (Manager) Draft refers to assignments generated by random serial dictatorship ([Abdulkadiroğlu and Sönmez, 1998](#)) led by workers (managers), while Match Quality Minimizer refers to matches generated from the Kuhn-Munkres algorithm that minimize the total firm objective (while still ensuring all workers are matched). 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 C.8 contains variations of this table and a regression version.

C.8.1 Worker and manager ranking of assignment, by matching algorithm, for only successful matches

In Tables 5 and C10 of the main text, we document that workers (managers) are matched to less-preferred managers (workers) in the firm-dictated match versus the DA-generated one. That table displays the average ranking across matching algorithms for all worker-manager pairs, including unmatched managers. In Table C11 below, we reproduce the same results but limit our sample to only successfully-matched worker-manager pairs. The results are in line with Tables 5 and C10.

Table C11: Manager Ranking of Assignment by Matching Algorithm (Matched Managers Only)

	Mean	Min	P25	Median	P75	Max
Managers-Propose DA	1.19 (0.05)	1.00 (0.00)	1.00 (0.00)	1.00 (0.00)	1.00 (0.00)	6.36 (0.62)
Workers-Propose DA	1.19 (0.04)	1.00 (0.00)	1.00 (0.00)	1.00 (0.00)	1.00 (0.00)	6.36 (0.62)
Worker Draft	1.32 (0.07)	1.00 (0.00)	1.00 (0.00)	1.00 (0.00)	1.00 (0.25)	7.44 (1.17)
Manager Draft	1.63 (0.08)	1.00 (0.00)	1.00 (0.00)	1.00 (0.00)	2.00 (0.06)	10.26 (1.95)
Firm-Dictated	2.99 (0.35)	1.00 (0.00)	2.00 (0.03)	2.00 (0.42)	4.00 (0.65)	12.00 (1.36)
Random Assignment	3.00 (0.36)	1.00 (0.00)	2.00 (0.00)	2.00 (0.47)	3.78 (0.65)	11.54 (1.36)
Match Quality Minimizer	3.19 (0.41)	1.00 (0.00)	2.00 (0.04)	3.00 (0.50)	4.00 (0.68)	11.00 (1.01)

Notes: This table displays the average manager ranking for matches generated by various algorithms. Worker (Manager) Draft refers to assignments generated by random serial dictatorship [Abdulkadiroğlu and Sönmez \(1998\)](#) led by workers (managers), while Match Quality Minimizer refers to matches generated from the Kuhn-Munkres algorithm that minimize the total firm objective (while still ensuring all workers are matched). Standard errors of the mean were bootstrapped by the quarter of the market, using one thousand bootstrapped samples per matching strategy. Because there were more managers on the market than workers, some managers are unmatched. This table includes only complete manager worker/pairs (i.e., dropping managers who were not paired to workers). All matching algorithms left the same number of managers unmatched, but differed in their composition. A similar table including unmatched managers (containing a score of zero) is accessible in [Tables 5 and C10](#).

C.8.2 Percentile of worker and manager ranking of assignment, by matching algorithms

In Tables 5 and C10 of the main text, we document that workers (managers) are matched to less-preferred managers (workers) in the firm-dictated match versus the DA-generated one. In Table C11 below, we reproduce the same results but using the percentile rank of each match. For each worker (manager), we calculate the percentile rank as $PR = \frac{CF - (.5 * F)}{N} * 100$, where CF is the count of all rankings less than or equal to the rank of interest, F is the frequency of the rank,¹⁰ and N is the number of managers (workers) available for the worker (manager). The results are also in line with Tables 5 and C10.

Table C12: Worker Percentile Ranking of Assignment by Matching Algorithm

	Mean	Min	P25	Median	P75	Max
Workers-Propose DA	2.72 (0.82)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	58.93 (4.28)
Managers-Propose DA	2.73 (0.82)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	58.93 (4.28)
Worker Draft	2.79 (0.67)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	51.14 (0.53)
Manager Draft	14.84 (1.59)	0.00 (0.00)	0.00 (0.00)	0.00 (0.37)	48.27 (13.13)	61.62 (3.04)
Command-and-Control	44.17 (1.14)	0.00 (0.00)	48.84 (1.61)	49.21 (0.06)	50.80 (0.61)	65.52 (3.85)
Random Assignment	44.57 (0.65)	0.00 (0.00)	48.83 (0.26)	49.21 (0.08)	50.80 (0.59)	65.64 (2.54)
Match Quality Minimizer	48.45 (0.37)	0.00 (0.00)	49.17 (0.19)	49.32 (0.28)	51.16 (0.58)	64.29 (2.84)

Notes: This table displays the average worker percentile ranking for matches generated by various algorithms. Worker (Manager) Draft refers to assignments generated by random serial dictatorship [Abdulkadiroğlu and Sönmez \(1998\)](#) led by workers (managers), while Match Quality Minimizer refers to matches generated from the Kuhn-Munkres algorithm that minimize the total firm objective (while still ensuring all workers are matched). Standard errors of the mean were bootstrapped by the quarter of the market, using one thousand bootstrapped samples per matching strategy.

¹⁰Some rankings appear multiple times given workers left some managers unranked, and we input these as tied.

Table C13: **Manager Percentile Ranking of Assignment by Matching Algorithm**

	Mean	Min	P25	Median	P75	Max
Managers-Propose DA	17.06 (2.39)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	2.00 (18.19)	99.06 (0.27)
Workers-Propose DA	17.07 (2.39)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	2.00 (18.15)	99.06 (0.27)
Worker Draft	17.45 (2.46)	0.00 (0.00)	0.00 (0.00)	0.00 (0.17)	5.11 (16.90)	99.06 (0.30)
Manager Draft	26.41 (2.43)	0.00 (0.00)	0.00 (0.00)	0.00 (0.81)	49.96 (0.42)	99.06 (0.38)
Command-and-Control	50.85 (0.79)	0.00 (0.00)	48.92 (0.29)	50.00 (0.42)	51.90 (0.84)	98.65 (0.38)
Random Assignment	51.16 (0.86)	0.00 (0.00)	48.82 (0.23)	49.97 (0.42)	51.92 (0.92)	98.62 (0.25)
Match Quality Minimizer	54.50 (0.99)	0.00 (0.00)	49.05 (0.18)	50.00 (0.35)	52.63 (1.12)	99.03 (0.38)

Notes: This table displays the average worker and manager percentile ranking for matches generated by various algorithms. Worker (Manager) Draft refers to assignments generated by random serial dictatorship [Abdulkadiroğlu and Sönmez \(1998\)](#) led by workers (managers), while Match Quality Minimizer refers to matches generated from the Kuhn-Munkres algorithm that minimize the total firm objective (while still ensuring all workers are matched). Standard errors of the mean were bootstrapped by the quarter of the market, using one thousand bootstrapped samples per matching strategy.

Table C14: **Worker losses from firm-dictated match**

	Manager unranked by worker	Worker ranking of manager	Simulated ranking of manager	Worker utility, z-score
Firm-Dictated	0.82*** (0.031)	2.16*** (0.44)	23.8*** (2.64)	-0.47*** (0.087)
R^2	0.762	0.546	0.516	0.470
Observations	64200	64200	64200	64200
DA mean	0.062	1.172	2.578	0.704
Fixed effect	Worker	Worker	Worker	Worker

Notes: This table examines the effect of using the firm-dictated matching system on worker preferences. The table displays the results of regression of four measures of worker preferences on an indicator for the firm-dictated match, with worker fixed effects. The regression drops worker-manager pairs obtained from random matching, so the comparison group is matching from deferred acceptance. Each regression uses robust standard errors clustered at the simulation-submarket level. Some workers did not rank managers and vice-versa. In column 2, we impute these as workers being indifferent among these choices (if a worker submitted N rankings, we impute the non-ranked managers as the worker's N+1 choice. In column 3, we randomly assign preferences over unranked choices. Column 5 uses a standardized utility measure from a rank-ordered logit, and drops some workers/managers who only submitted one ranking.

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

Table C15: **Manager losses from firm-dictated match**


	Worker unranked by manager	Manager ranking of worker	Simulated ranking of worker	Manager utility, z-score
Firm-Dictated	0.68*** (0.046)	0.93*** (0.19)	16.0*** (2.47)	-0.097** (0.030)
R^2	0.749	0.428	0.552	0.463
Observations	77550	77550	77550	77250
DA mean	0.225	8.202	9.357	-0.698
Fixed effect	Manager	Manager	Manager	Manager

Notes: This table examines the effect of using the firm-dictated matching system on worker and manager preferences. The table displays the results of regression of four measures of manager preferences on an indicator for the firm-dictated match, with manager fixed effects. The regression drops worker-manager pairs obtained from random matching, so the comparison group is matching from deferred acceptance. Each regression uses robust standard errors clustered at the simulation-submarket level. Some workers did not rank managers and vice-versa. In column 2, we impute these as workers being indifferent among these choices (if a worker submitted N rankings, we impute the non-ranked managers as the worker's N+1 choice. In column 3, we randomly assign preferences over unranked choices. Column 5 uses a standardized utility measure from a rank-ordered logit, and drops some workers/managers who only submitted one ranking.

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

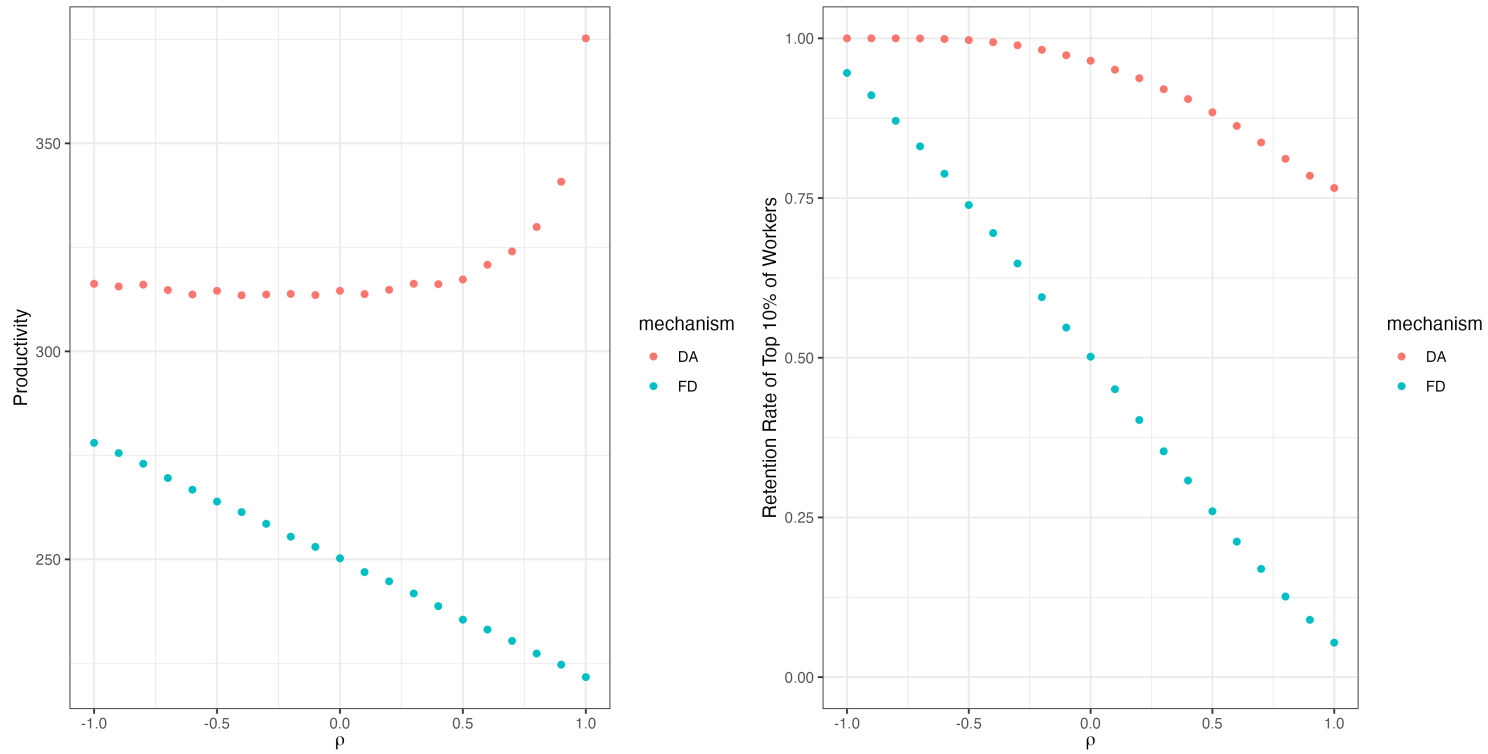
C.9 Additional Figures and Tables

Figure C3: DA Retention Benefits with Preference Correlation




Notes: This figure shows average retention outcomes from 1,000 draws of a 50×50 matching market where underlying utilities over matches are multivariate normally distributed with mean zero, unit variance, and eleven different correlations. The x-axis shows the correlation in the proposing side's preferences. The receiving side's preferences are uncorrelated. The different colors of lines show different unconditional retention rates. The left panel shows the retention rate of the proposing side, and the right panel shows the retention rate of the receiving side.

Figure C4: Productivity and Retention with Correlated Outside Options



Notes: This figure shows average outcomes across 10,000 simulations of a 100×100 matching market where divisions have vertical preferences over workers (based on each worker's productivity which is identical across all matches) and worker productivity is correlated with outside options. The left panel shows average productivity after accounting for retention. The right panel shows average retention of the 10 most productive workers in each simulation. The x-axis shows the correlation between each worker's productivity and outside option. Worker's preferences over matches are independent and identically distributed standard normally.

Figure C5: Empirical \mathcal{V}



Notes: This figure displays the histogram of the total output across a sample of 10K random assignments of workers and managers in our sample. On the right is the CDF of this distribution. Because of the large number of possible assignments (over 10^{35} in our smallest market alone), we visualize only a random sample of 10K possible matches in this figure. Although the density approaches zero near the edges of the figure, the full PDF extends broadly to the left and right.

Table C16: **Preference Alignment: Rank-Ordered Logit**

	Worker ranked manager #1		Manager ranked worker #1			
Objective Score (v_{ij}), σ	0.013*** (0.0027)	0.0093*** (0.0027)	0.017*** (0.0026)	0.014*** (0.0026)		
Firm-Dictated Assignment, σ		0.014*** (0.0032)	0.012*** (0.0032)	0.014*** (0.0032)	0.012*** (0.0032)	
R^2	0.000	0.000	0.000	0.000	0.000	0.000
Observations	23361	23361	23361	23361	23361	23361
Pseudo-R2	0.000	0.000	0.000	0.000	0.000	0.000

Notes: This table examines the alignment between worker/manager preferences and the firm's preferences. The units of observation are a worker \times manager pair. Columns 1–3 display the results of rank-ordered logit regression predicting all workers' rankings 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 the manager's rankings. Table 6 shows similar results using OLS specifications of which pairs were ranked #1 by either side. 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.


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

C.10 Distribution of worker and manager quality

Some workers and managers have high (or low) match quality across all possible match partners, on average. Although there is wide variability in agents' best and worst match (above), we also find strong individual fixed effects. Regressions predicting the match quality score for a match using workers FEs alone have an adjusted R^2 of 0.62. The adjusted R^2 is 0.58 for manager FEs alone and is 0.85 using both FEs together.

Figure C6 displays histograms of worker and manager quality. We take the average firm match quality score for each worker across all managers and de-mean this measure.

Figure C6: Histograms of worker and manager quality (de-meaned)




Notes: This figure displays a histogram of worker and manager quality. We take the average firm match quality score for each worker (manager) across all managers, de-mean this measure, and plot the distribution across all workers.

C.11 Distribution of worker specialization

In Definition 6 of our theory section, we say that a workforce is unspecialized if the workers are equally valuable in all matches. We can study this empirically by examining how much a worker's match quality score changes with their assigned manager. We take the difference between the quality of each worker's best match and worst match, and then divide this difference by the standard deviation of the firm match quality score across all matches. Figure C7 displays a histogram of this measure.

Figure C7: Histogram of worker specialization



Notes: This figure displays a histogram of worker specialization. We take the difference between the quality of each worker’s best match and worst match, and then divide this difference by the standard deviation of the firm match quality score across all matches.

C.12 Uncertainty about Firm Payoffs


In Figure C8, we present a series of simulations in which we add noise to all objective score values (v_{ij}). We needed to add noise to each objective score while respecting the upper and lower bounds created by the firm. To achieve this, we transformed each objective score variable into an unconstrained value; applied the noise, and then transformed them back into the original scale. Because our data was constrained to fall within $[0,1]$, we used the logit transformation.

Specifically, we first produced an unconstrained value of $x_{ij} = \text{logit}(v_{ij}) = \log(v_{ij}/(1 - v_{ij}))$. In cases where v_{ij} equaled zero or one, we added or subtracted an ϵ of 0.0001 to avoid infinities (our results were not sensitive to the exact choice of epsilons). Then, we added random Gaussian noise to create $y_{ij} = x_{ij} + \eta$, where $\eta \sim \mathcal{N}(0, \sigma^2)$. Finally, we transformed y_{ij} back into the $[0,1]$ scale by applying the inverse logit transformation to produce a set of $\tilde{v}_{ij} = \exp(y_{ij}/(1 + \exp(y_{ij})))$. The resulting \tilde{v}_{ij} s are garblings of v_{ij} .

We then use these garbled values as “ground truth” – and use the originals in our data as noisy observations of it. Our simulation captures the idea that the original v_{ij} s (used to calculate the assignments) were potentially noisy observations of an underlying true performance value. For ease of interpretation in the graph below, we plot the difference between the average level of output between firm-dictated assignment and DA. In our first set of simulations on the left side, we produced very low σ garblings that were highly correlated with the original v_{ij} s. We gradually increased the σ until it produced

correlations of zero.

Figure C8: Noise in Payoffs and Performance of Firm-Dictated Matching



Notes: The figure above shows the results of a series of simulations in which we add noise to all objective score values (v_{ij}). We then use these garbled values as “ground truth” – and use the originals in our data are noisy observations of it. We examine low levels of noise (i.e., garbled values correlated ≈ 1 with the original data) up to high levels (correlation of 0). The results show how noisily firm objectives must be known in order to enjoy varying degrees of benefits from firm-dictated assignment.

C.13 What does the workforce value?

If the workforce does not value the principal’s objectives, what do they value? We answer this question using data from two sources.

Survey Data. The first is an anonymous survey administered to the initial wave of workers who went through the market. Workers were asked how important different factors were in their rankings of opportunities. Table C17 summarizes the responses. Our survey results suggest several specific objectives that workers prioritize many objectives besides the firm’s production output.

On average workers' most important consideration was the "role description" containing what tasks the worker would be doing on a daily basis. The second most important consideration was the identity of the manager. The third most important was which product line the worker would be assigned.

"Skill fit and qualifications" – one of the main goals in the executives rankings – was tied for second-to-last in the workers' rankings with "office building and location" within the company campus. These locations were often correlated with the convenience of workplace amenities such as breakrooms, lunch and exercise/gym locations. Finally, workers ranked *role duration* as the least important factor. One possible reason is that the duration of roles did not vary much. 95% of roles had a duration between nine and fifteen months (with a median of twelve months).

Table C17: Survey Results: Important Factors in Ranking of Opportunities

<i>Job Characteristic</i>	<i>Mean Rating</i>	<i>SD</i>
Role Description	3.55	0.78
Manager	3.33	0.84
Product Line	3.27	0.95
Building & Location	3.11	1.02
Skill Fit	3.11	1.02
Duration	2.77	1.11

Notes: Workers were asked to label how important each job characteristic was in their rankings of possible matches. The survey was answered by $N = 18$ workers in the first round. The options were "unimportant," "neutral," "important" and "very important." We have coded these numerically as 1 through 4.

These survey results shed some light on what workers care about (besides maximizing firm output).

Profile Data. The second piece of evidence comes from workers' profile data. The matching website features a feature in which workers' skills and experience is described in a structured way around ≈ 200 skills. Within this framework, workers can label which skills they would like to grow in future assignments. In order to connect our work to prior literature on skill development (Pigou, 1912; Becker, 1964), we manually coded each skill as either general or firm-specific. $\approx 85\%$ of the skills were general (i.e., not firm-specific).¹¹

We summarize our findings in Table C18. Our first two columns results suggest that growth goals are private multidimensional information to the worker. The worker's growth goals cannot be predicted by worker or skill fixed effects (Column 1), or the

¹¹85% may be an underestimate. Several skills involved the effective use of the firm's products – which are also used by clients in the broader industry. Nonetheless we categorized these as "firm-specific." In addition, we have coded as "firm-specific" some skills that are required at other companies ("bug management"), but where the firm's internal processes may be idiosyncratic.

interactions between these two fixed effects and the performance level (performing or mastery) of the goals (Column 2).

With all controls included (Column 2), the R^2 is still 0.59. In the remaining columns of Table C18, we show that workers tend to label skills as “interested in growth” where they are not currently at a performing or mastery level. As a result, workers who prioritize their own growth are in tension with the firm.

Table C18: Skill Level and Desire for Growth

			Skill Labeled as Growth Desired		
Assessed Performing or Mastery			-.086*** (.021)	-.075*** (.015)	-.075*** (.017)
Worker FEs	Y	Y		Y	Y
Skill FEs	Y	Y			Y
All Interactions		Y			
Mean Dep. Var	.22	.22	.22	.22	.22
R^2	.48	.59	.0085	.43	.43
Observations	207,955	207,944	207,965	207,955	207,955

Notes: This table predicts whether a worker will label a particular skill as “growth desired,” as a function of whether this skill was previously assessed as “performing or mastery” in prior performance.

Our next set of results relates the desire for growth to the workers’ preferences over jobs. For each worker \times job pair, we count the number of the workers’s growth skills that appear in the job’s list of required or preferred skills. In Table C19, we show regressions where the dependent variable is the rankings. We examine workers’ entire preferences rankings using a rank-ordered logit approach, as well as the workers’ top choice.

In both specifications, we find that workers place a positive weight on jobs that allow them to grow the skills they desire to expand – even after controlling for the degree of broad and narrow alignment in these jobs. Together, these two tables show how workers’ desire for growth on privately-known skills drives differences between the firm’s objectives and the worker’s.

Table C19: **Workers 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*** (0.0037)	0.0049** (0.0020)
Objective Score (v_{ij}) σ	0.0094*** (0.0027)	0.012*** (0.0026)
Firm-Dictated Assignment, σ	0.012*** (0.0033)	0.0071*** (0.0021)
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 \times manager pair. Column 1 displays the results of rank-ordered logit regression predicting all workers' rankings. Column 2 shows similar results using OLS specifications of which pairs were ranked #1 by the worker. OLS specifications include worker and manager fixed effects. Robust standard errors clustered at the worker are included.

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

References

- Abdulkadiroğlu, Atila and Tayfun Sönmez**, "Random serial dictatorship and the core," *Econometrica*, 1998, 66 (3), 689–701.
- Agarwal, Nikhil and William Diamond**, "Latent indices in assortative matching models," *Quantitative Economics*, 2017, 8, 685–728.
- 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.
- Diamond, Peter and Joseph Stiglitz**, "Increases in risk and in risk aversion," *Journal of Economic Theory*, 1974, 8 (3), 337–360.
- Gale, David and Lloyd S. Shapley**, "College Admissions and the Stability of Marriage," *The American Mathematical Monthly*, 1962, 69 (1), 9–15.
- Kuhn, Harold W**, "The Hungarian method for the assignment problem," *Naval research logistics quarterly*, 1955, 2 (1-2), 83–97.
- Pigou, Arthur Cecil**, *Wealth and welfare*, Macmillan and Company, limited, 1912.

- Pittel, Boris**, “The average number of stable matchings,” *SIAM Journal on Discrete Mathematics*, 1989, 2 (4), 530–549.
- , “On Likely Solutions of a Stable Marriage Problem,” *The Annals of Applied Probability*, 1992, May, 358–401.
- Roth, Alvin E**, “Deferred acceptance algorithms: History, theory, practice, and open questions,” *international Journal of game Theory*, 2008, 36, 537–569.
- Roth, Alvin E. and Marilda Sotomayor**, *Two-Sided Matching: A Study in Game-Theoretic Modeling and Analysis*, Cambridge, UK: Cambridge University Press, 1990.