Delegated Recruitment and Statistical Discrimination*

Stepan Aleksenko† Jacob Kohlhepp‡

May 2, 2024

Abstract

We study how delegated recruitment shapes talent selection. Firms typically pay recruiters via refund contracts, which specify a payment upon the hire of a suggested candidate and a refund if a candidate is hired but terminated during an initial period of employment. We develop a model where refund contracts naturally arise and show that delegation leads to statistical discrimination, where the recruiter favors candidates with more precise productivity information. This is misaligned with direct hiring, where the firm has option value from uncertain candidates. Under tractable parametric assumptions, we characterize the unique equilibrium in which candidates with lower expected productivity but more informative signals ("safe bets") are hired over candidates with higher expected productivity but less informative signals ("diamonds in the rough").

Keywords: delegation, contracts, recruiters, screening, discrimination

JEL Codes: D82, D83, J24

Declarations of interest: none

*We are grateful for comments and suggestions from Alex Bloedel, Simon Board, Moritz Meyer-ter-Vehn, Tara Sinclair, and participants in the UCLA Theory Proseminar, UNC Theory Seminar, the SEA Annual Meeting, and SOLE Annual Conference. This project benefited from the support of the Institute for Humane Studies. We thank Hanqing Ye for excellent research assistance.

†Department of Economics, UCLA. Contact: saleksenko@ucla.edu.

‡Corresponding author. Department of Economics, UNC Chapel Hill. Contact: kohlhepp@unc.edu.

1 Introduction

The use of delegated recruitment has increased in the past few decades. Between 1991 and 2022, the fraction of American workers who found their job through a recruiter or headhunter rose from 4.9% to 14.3%, at the expense of direct applications (Black, Hasan, and Koning 2022). How does delegation impact the choice of candidates hired?

To answer this question, we develop a model where a firm (the principal) wishes to fill a position. A recruiter (the agent) can choose to suggest a candidate based on imperfect private information about the candidate's productivity. Candidates differ in their underlying productivity and the amount of information available about their productivity. If a candidate is suggested by a recruiter, the firm pays a hiring cost, observes the candidate's true productivity, and then chooses whether to terminate or retain the candidate. Given that employment is observable but the candidate's productivity and the recruiter's signals are not, all feasible contracts can be expressed as refund contracts, typical in the recruiting industry. Such contracts consist of a payment when a candidate suggested by the recruiter is hired by the firm, and a refund when a candidate is hired by the firm but terminated during an initial period of employment.

We demonstrate that refund contracts generate misalignment over whom to hire. Because refund contracts penalize recruiters for terminated candidates but do not reward recruiters for suggesting outstanding performers ("stars"), the recruiter prefers candidates whose productivity is likely to exceed the threshold. This leads to statistical discrimination in favor of more information when the contract features strong screening incentives: the recruiter requires a lower expected productivity threshold to suggest candidates whose resumes the recruiter understands better.

In contrast, in a first-best benchmark with direct recruitment, the firm has the option value of trying candidates with large residual uncertainty because of the firm's ability to terminate poor performers and retain stars. This leads to group statistical discrimination in favor of less information: the firm requires a lower expected productivity threshold to suggest candidates whose resumes the firm understands less.

We further show that misalignment occurs when the firm designs an optimal contract under a tractable Pareto-uniform information structure. The unique equilibrium contract features strong screening incentives: the recruiter suggests only candidates who will certainly not be terminated. As a result, delegation reverses the direction of statistical discrimination compared to the first-best benchmark. Furthermore, low-expected-productivity candidates about whom the recruiter is better informed ("safe bets") are inefficiently hired at the expense of high-expected-productivity candidates about whom the recruiter is less informed

("diamonds in the rough"). Thus, delegation via refund contracts systematically disadvantages groups of candidates. All groups with less informative signals than a cutoff are hired with a lower probability than in the first-best benchmark, while all groups with more informative signals than a cutoff are hired with a higher probability than in the first-best benchmark.

Information heterogeneity across candidates is a key driver of inefficiency in equilibrium. When the recruiter observes the same amount of information about all candidates, the first-best benchmark is achieved. Information heterogeneity can reflect both differences in a candidate's profile (e.g., one candidate may have worked at a well-known company or may have obtained a college degree from a well-known university, while another may have not) and differences in a recruiter's ability to interpret a candidate's profile (e.g., a recruiter may be better at interpreting the resumes of candidates that share the recruiter's race, gender, or socioeconomic backgrounds).[1] Under this interpretation, our results imply recruitment delegation favors the candidates who can afford to acquire precise but costly productivity signals and those who share demographic characteristics with the recruiter.

By considering extensions of the baseline model, we show that the qualitative results derived under the parametric information structure without commitment are robust. When the firm has a third payment it can use to extract all surplus from the agent, profit weakly increases but the first-best benchmark is not achieved even when information heterogeneity is small. When the recruiter can search for other candidates, the contract additionally serves to encourage unobserved effort, but different types of statistical discrimination between the first-best benchmark and equilibrium remain. When the firm can commit to a termination threshold, full extraction occurs and misalignment persists due to the fundamental tension between first- and second-stage screening.

Using a series of comparative statics in productivity and signal distributions, we show how delegation induces various spillovers across groups of candidates in equilibrium. Better information about some candidates makes screening unambiguously more valuable and therefore increases the hiring bar and decreases the probability of being hired for the rest of the candidates. However, better productivity for some candidates can either increase or decrease the hiring bar, depending on the quality of the information available for those candidates.

1. For an example of racial homophily in hiring, see Giuliano, Levine, and Leonard (2009).

1.1 Literature

Our work relates to several strands of literature. Similar to many papers in the literature on delegation to an expert, our model features an agent with an information advantage making decisions for a principal. However, unlike in many papers in this literature (e.g., Frankel 2014, Kundu and Nilssen 2020, Szalay 2005), in our paper the principal cannot constrain the action set of the agent directly and must exert control indirectly through monetary payments. Our paper shows refund contracts are a way to screen out recruiters with low-productivity candidates, similar to how debt contracts can be used to screen borrowers (Gale and Hellwig 1985). Additionally, the agent's bias stems from not internalizing the principal's hiring cost, similar in spirit to Che, Dessein, and Kartik (2013) where the agent does not internalize the firm's outside option. Notably, Che, Dessein, and Kartik (2013) find a bias towards "conditionally better looking" projects, similar to how we find a bias towards "safe-bet" applicants.

Our paper is also related to the literature on delegated information acquisition. As in Chade and Kovrijnykh 2016, Inderst and Ottaviani 2012, and Szalay 2009, the principal must consider how the agent will act when faced with varying amounts of information. However, while in these papers the amount of information acquired is endogenous and homogeneous, in ours it is exogenous but heterogeneous. Moreover, the focus of the principal in this literature is to encourage information acquisition, while in our paper it is to achieve the best screening behavior on average across differently informed agents.

Our model features an agent who is privately informed about the candidates' expected productivity and the amount of information available about it. As a result, our paper is related to the literature on multidimensional screening (Carroll 2017, Yang 2021). An essential aspect of our setting is that the agent is privately informed about the quality of information available. Even though productivity is the only object the firm cares about, the distribution of productivity across interim agent types differs in two dimensions. Additionally, the agent does not have intrinsic preferences over the dimensions being screened, as in problems of screening buyers. Instead, the way the agent engages with the contract offered by the firm is impacted by the agent's private type, and the firm exploits this in equilibrium to screen out recruiters with low-expected-productivity candidates. Another difference from this literature is that we do not allow the firm to offer a menu of contracts. This is because our model is application driven, and to our knowledge contract menus are not commonly observed in the recruiting industry.

Our paper differs from most past work in that the principal commits to a monetary contract, but this contract explicitly depends on an action taken by the principal after all information is revealed. There is a form of limited commitment built into the model that is

inspired by the specific application to recruiters. One exception is Levitt and Snyder (1997) where, as in our model, an agent observes a private signal about a project's success and the principal can cancel the project based on the agent's interim advice. A key difference is that in Levitt and Snyder (1997) the contract is accepted before information is observed. The authors find that the principal's ability to influence the final contractible outcome undermines incentives, and commitment can help rectify the situation.

Our paper also contributes to the literature that studies what occurs when the hiring decision is delegated. Frankel (2021) considers delegation to a hiring manager, where the trade-off is between using the manager's soft information and indulging the manager's bias towards soft information relative to hard information. Two key differences are that in Frankel (2021) the firm has access to hard information and uses it to limit the manager's actions directly. Cowgill and Perkowski (2020) propose a theoretical framework and study empirically how recruiters select applicants. They show in a two-sided audit that recruiters over-interview candidates from elite schools and big companies, a fact the authors interpret as evidence of a reputational effect. Our paper proposes another interpretation: even conditional on the same expected productivity of candidates, recruiters are biased towards candidates from elite schools or big companies because of the way refund contracts are structured.

After presenting the model, the paper proceeds as follows. In Section 3.1, we derive hiring thresholds in a first-best benchmark where the firm does not delegate but rather observes the candidate directly. In Section 3.2, we derive the equilibrium contract. In Section 3.3, we compare the set of hired candidates in the first-best benchmark and equilibrium and derive the main result that refund contracts induce misalignment and different types of statistical discrimination. In Section 4, we discuss the forces that lead to the main result. Finally, we analyze several comparative statics in Section 5, discuss the robustness of our results in Section 6, and conclude the paper in Section 7.

2 Model

Players and Actions. A risk-neutral firm wishes to hire one candidate, and a risk-neutral recruiter has one candidate with uncertain productivity (a). The firm proposes a contract to the recruiter. After receiving private information about the candidate's productivity, the recruiter either accepts the contract by suggesting the candidate or rejects the contract by not suggesting the candidate. If the candidate is suggested, the firm incurs a hiring cost (c), fully and privately observes the candidate's productivity, and then decides whether to retain

or terminate the candidate.[2] If the candidate is retained, the firm receives the candidate's productivity. Finally, all contract transfers are realized.

If the candidate is not suggested, both the firm and the recruiter receive their outside option. The outside option of the firm is 0, and that of the recruiter is $\bar{u} \geq 0$. The candidate is not a strategic player.

Candidates and Information. A candidate is characterized by a productivity a distributed according to a common prior with a CDF F_a and an information type (or candidate's group) $i \in \{1, ..., N\}$ with probabilities $\{p_1, ..., p_n\}$, which are drawn independently. Prior to deciding whether to suggest the candidate to the firm, the recruiter privately observes the candidate's i and a signal about productivity, $x \in \mathbb{R}^{\tau_i}$, distributed according to a CDF $G_i(\cdot|a)$ for the candidate with productivity a and information type i.[3] With a slight abuse of notation, we call G_i an information structure for the productivity signal that the recruiter observes for a candidate from group i. Throughout, we index candidate information types in descending order of their informativeness in the Blackwell sense, i.e., G_i is Blackwell more informative than G_j for any $i < j$. To make these concepts concrete, we provide two parametric examples below.

Normal example. The prior distribution of the candidate's productivity is normal: $a \sim N(\mu_a, \sigma_a^2)$; formally, F_a is the CDF of a normal random variable with mean μ_a and variance σ_a^2. For each information type, the recruiter observers a single productivity signal ($\tau_i = 1$), which is the sum of true productivity and independent normal noise: $x = a + \varepsilon$, where $\varepsilon|a, i \sim N(0, \sigma_i^2)$. Using our notation, $G_i(\cdot|a)$ is a CDF of a normal random variable with mean a and variance σ_i^2. Groups with lower i have higher signal precision ($\sigma_1^2 < \sigma_2^2 < ... < \sigma_N^2$). Conditional on observing signal x for a candidate with information type i, the recruiter's posterior belief about productivity is $N\big(\frac{\sigma_i^2}{\sigma_a^2\sigma_i^2} \cdot \mu_a + \frac{\sigma_a^2}{\sigma_a^2\sigma_i^2} \cdot x, \frac{\sigma_a^2\sigma_i^2}{\sigma_a^2+\sigma_i^2}\big)$.

Pareto example. The prior distribution of the candidate's productivity is Pareto: $a \sim Pareto(\bar{a}, k)$. The recruiter observes a productivity signal consisting of τ_i unidimensional signals: $x = (x_1, ...x_{\tau_i})$. Groups with lower i have more signals: $\tau_1 > \tau_2 > ... > \tau_N$. Conditional on productivity a, signals are drawn i.i.d. from a uniform distribution with minimum 0 and maximum a. Using our notation, $G_i(\cdot|a)$ is a CDF of a multivariate uniform distribution on $[0, a]^{\tau_i}$. Conditional on observing signal x for a candidate with information type i, the recruiter's posterior belief about productivity is $Pareto(\max\{\bar{a}, \{x_t\}_{t=1}^{\tau_i}\}, \tau_i + k)$.

Contracts and Payoffs. The productivity signal x is the recruiter's private information and is not contractible. The productivity a is privately observed only by the firm and is also

2. After suggestion but before paying c, the firm does not observe any additional information about the candidate.

3. The candidate's information type i (or group) is not payoff relevant and serves only as a part of information structure description and, hence, the productivity signal.

not contractible. Even though the information type i is sometimes observable by both parties, we will frequently interpret it as representing demographic characteristics (race, gender, age, etc.). Because in many countries these characteristics are illegal to contract on, we assume they are also not contractible. This leaves only two possible contractible outcomes after the contract is accepted: whether the candidate is eventually retained or terminated.

Given the space of contractible outcomes, we can express all contracts as *refund contracts*, which consist of a transfer from the firm to the recruiter if the candidate is suggested and retained ($\alpha \in \mathbb{R}$) and a refund from the recruiter to the firm if the candidate is terminated ($\beta \in \mathbb{R}$). Under these contracts, the ex post profit of the firm is

$$\pi = \mathbf{1}_{\{suggested\}} \cdot \Big(-c - \alpha + \mathbf{1}_{\{retained\}} \cdot a + \mathbf{1}_{\{terminated\}} \cdot \beta \Big),$$

and the ex post utility of the recruiter is

$$u = \mathbf{1}_{\{suggested\}} \cdot \Big(\alpha - \beta \cdot \mathbf{1}_{\{terminated\}} \Big).$$

In Section 2.1, we argue that these contracts are commonly observed in the recruiting industry. We specify that the recruiter chooses to accept the contract if indifferent.

We consider two versions of the model in this paper. **A.** In the commitment case, the firm can commit to a termination threshold γ (for all candidates regardless of their information type i), include it as a part of the contract, and then terminate the candidate if and only if the candidate's realized productivity a is below γ. The contracts cannot depend on the realized productivity in any other way apart from the termination rule.[4] **B.** In the no-commitment case, the firm cannot commit to the termination rule and decides whether to terminate the candidate in the last stage of the game. Thus, the firm prefers to terminate the candidate if and only if the candidate's productivity a is below the refund β regardless of the candidate's information type i (thus, the ex post compatible termination threshold in the no-commitment case is $\gamma = \beta$).

Equilibrium. The equilibrium concept we use is weak Perfect Bayesian Equilibrium with the assumption that the recruiter has passive beliefs about the productivity of the candidate in the contract offered by the firm. We require passive beliefs in order to have well-defined beliefs under off-path contracts.[5] We assume that when indifferent, the recruiter suggests

4. If the contract could freely depend on the productivity, the firm would prefer to lie about its realization conditional on retaining or terminating the candidate, in order to increase its payoff.

5. Without passive beliefs, the recruiter does not have to be Bayesian when interpreting the signals under off-path contracts because those information sets are of zero probability. Thus, for any realization of signals for any off-path contract, the recruiter can be extremely pessimistic about the candidate's productivity and

the candidate. Finally, we assume that there exists a feasible interior suggestion strategy[6] which delivers strictly greater total surplus than always or never suggesting a candidate.

2.1 Model Comments

Refund Contracts. Refund contracts are a natural formalization of the contract space in our model. They closely resemble what is called guarantee contracts, which are the main form of contract used by external recruiters, according to evidence from qualitative interviews, industry materials, and surveys. We interviewed one mid-career headhunter and one early-career recruiter. Both stated they were paid if a candidate they suggested was placed, but they had to provide a refund or a free replacement if the candidate left the firm in the first 90 days of employment. This refund is called in the industry a "guarantee." The early-career recruiter confirmed that the refund was given for any reason, including termination of the candidate by the company. The relevant sections of the interview transcript are provided in Appendix Section 8.6.

Recruiters that use this structure of compensation are called contingent recruiters, and they represent the majority of recruiters, according to estimates (Finlay and Coverdill 2007). The contingent compensation structure with a guarantee discussed by the recruiters we interviewed is also mentioned in a variety of sources, including a report on recruiting practices in the hospitality industry (Dingman 1993), how-to books about starting an executive recruiting firm (Press 2007, Perry and Haluska 2017), a guide for lawyers working with headhunters (Steinberg and Machlowitz 1989), a guide for managing financial service companies (Arslanian 2016), and an academic article (Florea 2014). Further, the American Staffing Association provides a "Model Recruiting Agreement" which includes sample language for refund and replacement guarantees (Association 2014).

A survey of recruiters by Top Echelon found that 96% of recruiters offered some form of guarantee. The most popular guarantee time frame was 90 days, consistent with our interviews. When asked about the form of the refund, 61% responded that they offered a replacement, and 26% responded that they offered money back (Deutsch 2019).[7]

Information Types. The information type of the candidate may in some cases be observed by the firm, but it is unnatural and often illegal to condition on it in a contract.

suggest no one. Therefore, any contract can be an equilibrium one under wPBE. To fix that, we use passive beliefs and force the recruiter to make a Bayesian inference from the productivity signals in every information set (including those off-path).

6. Feasible by the firm via refund contracts, and "interior" meaning the probability of suggestion is not 0 or 1.

7. Of the remaining 13%, 11% responded that they offered a guarantee that did not fit into the two aforementioned categories, and 2% gave no response.

For example, the information type of a candidate may be determined by the candidate's demographics (age, gender, race, etc.). In the U.S. it is illegal to recruit based on such characteristics, much less write them explicitly in a contract.[8] In our model, the information type by itself is uninformative about the candidate's productivity. Alternatively, the information heterogeneity can also come from differences in the candidates' experiences or certifications, which might be unobserved by the firm and therefore cannot be included in the contract.

Firm and Recruiter. The hiring cost can be interpreted as the cost to the firm of interviewing the candidate[9] or the cost of employing the candidate for a probationary period. Within the model, the hiring cost operates as the cost to the firm of fully learning the candidate's productivity. This is an important ingredient in the model because it makes the recruiter's private information valuable to the firm. The firm's main goal is to use the recruiter's private information to screen candidates. Because the recruiter does not bear any intrinsic cost of suggesting a candidate, the recruiter has a natural tendency not to screen candidates and instead suggest everyone. For this reason, the contract must be designed to reduce this tendency.

It may appear that the only reason the firm uses a recruiter is that the recruiter has a candidate and the firm does not. However, by setting the payment from suggestion to be $\bar{u}$ and setting $\beta = 0$, the firm can always design a contract where the recruiter suggests everyone (i.e., no screening). Thus, if the firm had an outside option of paying the hiring cost c to learn the productivity of a random candidate or using a recruiter, it will always weakly prefer to use the recruiter. Recruiters are useful to the firm because of their information advantage prior to hiring, not because they have a monopoly on candidates.

3 Nonparametric Analysis

In this section, we define notions of statistical discrimination and show that discrimination goes in opposite directions in the first-best benchmark and equilibrium under a wide set of contracts. We begin by analyzing the space of posteriors about the candidate. Second, we define the first-best benchmark and equilibrium and derive the conditions under which the candidate is suggested. Third, we introduce directed notions of individual and group statistical discrimination and provide conditions under which the direction of statistical discrimination is different in the first-best benchmark and equilibrium. We finally show

8. Per EEOC (2023): "It is also illegal for an employer to recruit new employees in a way that discriminates against them because of their race, color, religion, sex (including gender identity, sexual orientation, and pregnancy), national origin, age (40 or older), disability or genetic information."

9. For example, the cost of having current employees conduct on-site interviews.

what forces distort the equilibrium outcome relative to the first-best benchmark.

3.1 Candidate Posteriors

Characterizing the recruiter's (or firm's) optimal decisions for each possible recruiter's information set is the same as characterizing the optimal suggestion decisions over the space of the candidate's group and productivity signal realizations. We start by parameterizing this space.

For any information type i and any realized signal x, we can derive the posterior productivity distribution $F_{i,x}(a)$. We assume that each group's information structure is such that no two signal realizations have the same posterior μ. This assumption allows us to uniquely map every posterior distribution of productivity $F_{i,x}$ to its mean μ and the information type of the candidate i, and therefore to parameterize the set of all candidates' productivity posterior distributions by i and μ ($F_{i,x} \leftrightarrow F_{i,\mu}$). Further, we describe the candidates upon observing their productivity signal and group by (i, μ).

To define intuitive risk attitudes for the firm and the recruiter, we require that the space of candidates satisfy the following assumption.

Assumption

(a) $F_{i,\mu_1}(a)$ *first-order stochastically dominates* $F_{i,\mu_2}(a)$ *for any* $(\mu_1 > \mu_2, i)$,

(b) $F_{i,\mu}(a)$ *second-order stochastically dominates* $F_{j,\mu}(a)$ *for any* $(\mu, i < j)$, *and*

(c) $\forall \mu, i < j \ \exists \ a^* \geq F_{j,\mu}^{-1}(1/2), \ s.t. \ \forall \ a < (>) a^* : \quad F_{i,\mu}(a) \leq (\geq) F_{j,\mu}(a)$.

Part (a) of this assumption means that there is an unambiguous order of signal realizations within an information type, where a "better" signal means that the candidate's posterior productivity distribution first-order dominates all posteriors with "worse" signals (conditional on the information type of the candidate). Parts (b) and (c) of this assumption require a stronger order of information structures than the Blackwell order. Besides requiring that the signal for a group j be a garbling of the signal for a group $i < j$, it requires that each posterior $F_{j,\mu}$ be a single-mean-preserving spread of $F_{i,\mu}$ and have a fatter lower tail than $F_{i,\mu}$.

Many common prior distributions and information structures satisfy this assumption, including (1) normal prior distribution with normal signals of various precision *(parametric example A)*; (2) the Pareto prior distribution with various numbers of uniform signals on $[0, a]$ *(parametric example B)*; and (3) any information structure with posteriors that can be expressed as $\mu + \sigma_i \varepsilon$, where ε is a symmetric mean-zero random variable.

3.2 The First-Best, Equilibrium, and Suggestion Decisions

First-Best. We define the first-best benchmark as a hypothetical situation where the firm possesses the recruiter's private information about the candidate and makes the hiring decision directly based on this information.

In the first-best benchmark (shortened to "first-best"), there is no need for the firm to worry about the payments to the recruiter when making the termination decision. The firm terminates the candidate only if $a < 0$ because the hiring cost is already paid. The firm's ex post value from a candidate with known productivity a is $\max\{a, 0\}$. The firm's expected value from hiring a candidate with information type i and expected productivity μ is equal to $\mathbb{E}[\max\{a, 0\}|i, \mu]$. Thus the firm suggests (and hires) a candidate if and only if: $\mathbb{E}[\max\{a, 0\}|i, \mu] \geq c$.

The first-order stochastic order in μ implies monotonicity of posterior truncated expectations in μ. We denote the lowest-posterior-expectation candidate from group i whom the firm suggests in the first-best as $\mu^*(i) = \inf_m\{m|\mathbb{E}[\max\{a, 0\}|i, \mu = m] \geq c\}$. Then the firm suggests a candidate (i, μ) if and only if $\mu \geq \mu^*(i)$. We focus our attention on non-trivial cases where the firm prefers to screen out a positive share of candidates in the first-best.[10]

Equilibrium. We begin our analysis of equilibrium by considering the firm's decision to terminate or retain the candidate. In the commitment case, the firm terminates the candidate if and only if the realized productivity a is below the termination threshold γ. In the no-commitment case, because the firm cannot commit to a termination rule, the decision is made after the contract is designed, hiring costs are sunk, and the firm has fully learned the productivity of the candidate. Thus the firm retains the candidate if the realized productivity exceeds the refund (β) that could be obtained from termination: $a \geq \beta$. This is the firm's incentive-compatible retention rule (IC). Thus, the termination threshold is given by $\gamma = \beta$ in the no-commitment case. For now, we proceed with a general γ (not necessarily IC).

Anticipating the firm's termination or retention decision, the recruiter has an ex post value from suggesting a candidate with known productivity a that is $\alpha - \beta\mathbb{I}_{a<\gamma}$. The recruiter's expected value from suggesting a candidate with information type i and expected productivity μ is equal to $\alpha - \beta \cdot \Pr(a < \gamma|i, \mu)$. Thus the recruiter suggests (and the firm hires) a candidate if and only if the interim value exceeds the recruiter's outside option: $\alpha - \beta \cdot \Pr(a < \gamma|i, \mu) \geq \bar{u}$. The recruiter's decision to suggest a candidate in equilibrium is therefore fully determined by whether a given candidate's termination probability is above or below a threshold: $\Pr(a < \gamma|i, \mu) < p^*$, where $p^* = \frac{\alpha - \bar{u}}{\beta}$.

10. If the firm prefers not to screen anyone out in the first-best, then the solution is trivial and is implementable in equilibrium.

With these results established, and recalling that our assumptions imply monotonicity of $\Pr(a < \gamma | i, \mu)$ in μ, we denote the lowest-posterior-expectation candidate from group i whom the recruiter suggests under a refund contract with parameters (α, β, γ) as $\mu_{\alpha,\beta,\gamma}(i) = \inf_m\{m | \Pr(a < \gamma | i, \mu = m) < \frac{\alpha - \bar{u}}{\beta}\}$. Then the recruiter suggests a candidate (i, μ) if and only if $\mu \geq \mu_{\alpha,\beta,\gamma}(i)$. We focus on non-degenerate cases where the firm's optimal contract induces some screening in equilibrium.[11]

3.3 Individual and Group Statistical Discrimination

In this section, we define individual and group statistical discrimination and show how delegation to the recruiter can generate both. We demonstrate that while the firm is endogenously risk-loving in the first-best, refund contracts often make the recruiter endogenously risk-averse.

We define individual discrimination based on how two candidates with the same expected productivity (μ) but different information types are treated.

Definition 1 *A decision maker engages in individual statistical discrimination in favor of more (less) information for candidates with expected productivity μ and information types $i < j$ if the candidate with the more (less) informative type i (j) is always suggested whenever the candidate with the less (more) informative type j (i) is suggested.*

To make this definition concrete, suppose there are two candidates (label them 1 and 2) from different information types (say, 1 and 2) but identical posterior productivity means: $(1, \mu)$ and $(2, \mu)$. Because candidate 1 comes from a lower information type, we are more informed about candidate 1. Specifically, $F_{2,\mu}(a)$ is a single-mean-preserving spread of $F_{1,\mu}(a)$. If candidate 1 being suggested implies that candidate 2 is suggested but not the other way around, then there is individual statistical discrimination in favor of less information. Intuitively, candidate 1 gives a higher value to the decision maker than does candidate 2. The next result establishes that individual statistical discrimination is quite different in the first-best compared to delegation via refund contracts.

Proposition 1

1. *In the first-best, the firm always engages in individual statistical discrimination in favor of less information (for all μ, i, j).*
2. *In equilibrium, the recruiter engages in individual statistical discrimination in favor of more (less) information if $\gamma \leq (\geq) a^*$, where $a^* : \ F_{i,\mu}(a^*) = F_{j,\mu}(a^*)$.*

11. If the firm prefers not to screen anyone out in the equilibrium, then the solution is trivial and leads to the recruiter suggesting all candidates and the firm retaining those with productivity above zero.

Proof. The first-best result follows directly from the fact that, due to the firm's ability to terminate candidates, the firm's value from a candidate in the first-best is a convex function of realized productivity ($\max\{0, a\}$). The refund contract result follows from the fact that information types are single-mean-preserving spreads and under refund contracts the recruiter value from a candidate depends only on the probability that a candidate is below a threshold. This value is exactly the CDF, so which candidate the recruiter prefers depends only on the termination threshold relative to the single-crossing point a^*.■

Proposition 1 captures the key tension which drives misalignment. Absent delegation, the firm is endogenously risk-loving; it appreciates the option value of uncertain candidates because of its ability to terminate candidates. However, refund contracts frequently cause the recruiter to be endogenously risk-averse. Specifically, the recruiter receives no additional reward for suggesting candidates with outstanding productivity but is penalized for suggesting candidates whose productivity falls below a threshold.

Whether the recruiter statistically discriminates in one direction or the other about a particular pair of candidates depends on the termination threshold. When the termination threshold γ is low (relative to the single crossing-point for that particular pair), the recruiter statistically discriminates in favor of candidates that the recruiter understands better. When the firing threshold γ is high (relative to the single crossing-point for that particular pair) the recruiter statistically discriminates in favor of candidates the recruiter understands less. This situation happens precisely because the recruiter's goal is to minimize the probability of termination for the suggested candidate.

The relation between the termination threshold (γ) and the single-crossing point a^* determines the amount of screening in the second stage of the hiring process (after the candidate was suggested). If γ is high (low) relative to a^*, a candidate is more (less) likely to be terminated (screened out) after being suggested by the recruiter. With little second-stage screening, the recruiter prefers less noise and more certainty in the candidate's productivity to make sure that it likely exceeds the threshold γ. With a lot of second-stage screening and many candidates that are unlikely to make the cut, the recruiter prefers more noise in the candidate's productivity, to have at least some chance of a high realization exceeding the threshold γ.

This result about individual discrimination is incomplete. The misalignment between the firm and the recruiter over any particular candidate depends on the termination threshold γ relative to the single-crossing point a^*. When this threshold is low, there is misalignment; when it is high, there is not. Because the single-crossing point between any two candidates (a^*) depends on the expectation and information type of those candidates, it is unclear how this individual alignment or misalignment of discrimination, which differs across specific

pairs of candidates, translates into systemic differential treatment of groups of candidates.

To address these concerns, we introduce the notion of group discrimination. Recall that $\mu^*(i)$ is the lowest-posterior-expectation candidate from group i that is suggested in the first-best, while $\mu_{\alpha,\beta,\gamma}(i)$ is the lowest-posterior-expectation candidate from group i that is suggested by the recruiter under a refund contract with parameters (α, β, γ).

Definition 2 *A decision maker engages in group statistical discrimination in favor of more (less) information in the first-best if $\mu^*(i)$ is increasing (decreasing) in i, and in equilibrium under refund contract (α, β, γ) if $\mu_{\alpha,\beta,\gamma}(i)$ is increasing (decreasing) in i.*

Group statistical discrimination captures whether the "hiring bar" is different across groups of candidates. The next result characterizes when differential group statistical discrimination occurs across the first-best and equilibrium.

Proposition 2

1. *In the first-best, the firm always engages in group statistical discrimination in favor of less information.*
2. *In equilibrium, the recruiter engages in group statistical discrimination in favor of more information if $\frac{\alpha-\bar{u}}{\beta} < \frac{1}{2}$.*

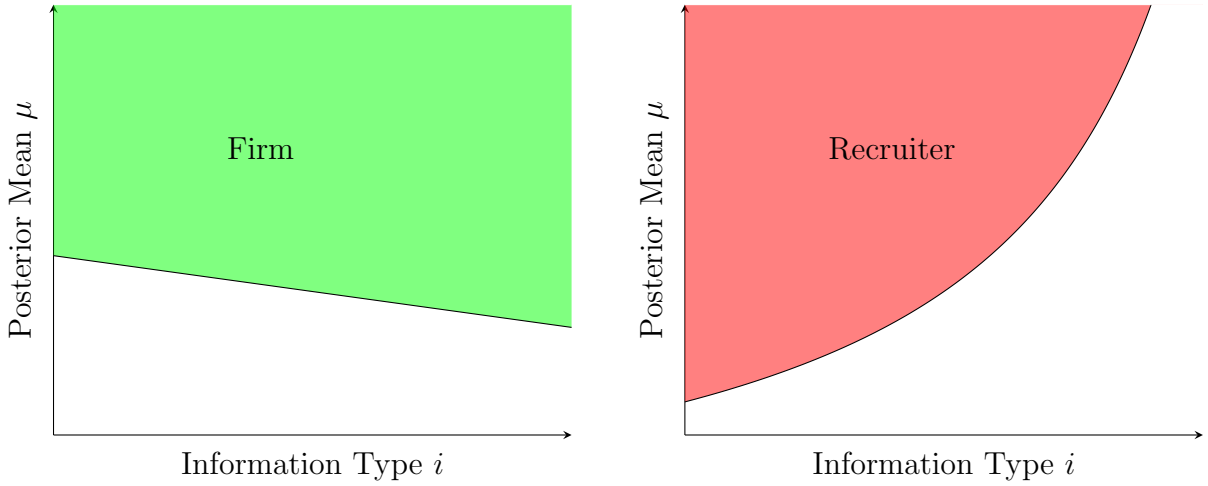
Proof. The fact that the first-best favors groups of candidates with less information follows directly from the previous proposition. For equilibrium, consider that the recruiter suggests any candidate when the recruiter's payoff from suggestion exceeds the outside option:

$$\alpha - \beta Pr(a < \gamma | i, \mu) \geq \bar{u} \leftrightarrow \frac{\alpha - \bar{u}}{\beta} \geq Pr(a < \gamma | i, \mu).$$

The threshold $\mu(i)$ is given by the points where this inequality holds with equality. When $\frac{\alpha-\bar{u}}{\beta} < \frac{1}{2}$, all candidates with less informative types will have a higher threshold $\mu(i) \leq \mu(j)$ due to single-crossing.■

Figure 1 illustrates Proposition 2 for $\frac{\alpha-\bar{u}}{\beta} < \frac{1}{2}$.

Figure 1: Candidate Types Hired in Equilibrium vs. First-Best



Proposition 2 shows that under a range of contracts, there are opposite types of statistical discrimination not only across individual pairs of candidates but also across groups of candidates. The firm sets a lower hiring bar for groups the firm understands less, while the recruiter sets a lower hiring bar for groups the recruiter understands more. Whenever this result holds, some groups will be disadvantaged in terms of hiring by delegated recruitment relative to direct recruitment.

The second part of Proposition 2 specifies that misalignment occurs under contracts where $\frac{\alpha-\bar{u}}{\beta} < \frac{1}{2}$. One way to interpret this condition is that the recruiter offers an economically significant refund. Notice that β is exactly the part of the contract that induces first-stage screening, i.e., screening in the suggestion decision. When $\beta = 0$, the recruiter suggests all candidates. When $\beta \to \infty$, the recruiter suggests only candidates that are retained with certainty. If β is high and therefore $\frac{\alpha-\bar{u}}{\beta}$ is low, many candidates are screened out in the first stage by the recruiter, and few of them are terminated after the probation period. Similar to individual discrimination, the key condition for misalignment with respect to group discrimination is little second-stage screening implied by the strong first-stage screening incentives.

An important observation from this condition is that the termination threshold does not directly determine whether group statistical discrimination in favor of more information occurs in equilibrium. However, the termination threshold will impact the way the firm designs the other contract parameters, so it will still play an indirect role. For this reason, in the next few sections, we analyze how the equilibrium contract is designed.

3.4 Contract Design in Equilibrium

The analysis up to this point treats the features of the contract (α, β, γ) as exogenous parameters. In this section, we discuss the key forces which determine how the contract is designed in equilibrium. We will also discuss the tensions between these forces that arise due to information heterogeneity within the candidate pool.

Corollary 1.

1. *$\mu_{\alpha,\beta,\gamma}(i)$ is decreasing in α (with or without commitment),*
2. *$\mu_{\alpha,\beta,\gamma}(i)$ is increasing in β (with or without commitment),*
3. *$\mu_{\alpha,\beta,\gamma}(i)$ is increasing in γ (with commitment).*

From the contract designer's (the firm's) perspective, the contract serves three purposes. We first consider the case where the firm can commit to a termination threshold (γ is a parameter of the contract). We conclude the section discussing the differences that arise without commitment ($\gamma = \beta$).

First-Stage Screening. The contract determines the set of candidates the recruiter suggests to the firm. The firm wishes to design the contract so that the equilibrium set of suggested candidates is close to the first-best set of suggested candidates. By increasing the suggestion payment α, the firm expands the set of candidates suggested across all information types. By increasing the refund β or the termination threshold γ, the firm shrinks the set of candidates suggested across all information types.

Replicating the first-best is often not possible because of the coarseness of refund contracts. They give the firm only three choice variables with which to control the suggestion or hiring region. This leads to a mechanical impossibility of achieving the first-best for more than three information types with commitment or more than two information types without commitment. Although we numerically verify that the first-best suggestion region is feasible with two information types, the firm does not choose to implement the first-best region in this case.

Second-Stage Screening. After hiring a candidate and paying the cost c, the firm fully learns the candidate's productivity and makes a termination decision (either mechanically with commitment or endogenously without). This can be viewed as a perfect but costly second stage of screening. The firm wishes to design the contract to minimize inefficient terminations: candidates with positive productivity ($a \geq 0$) that are hired at cost c and then terminated. Holding fixed the set of suggested candidates, reducing the positive termination threshold γ reduces these inefficient terminations.

Since the termination threshold γ affects both the hiring region (first-stage screening) and the termination threshold (second-stage screening), there is tension when choosing the optimal two-stage screening procedure. Raising the termination threshold improves first-stage screening but requires the firm to inefficiently terminate candidates after they are hired.

Surplus Extraction. Because the recruiter selects into the contract after observing signals about the candidate, the recruiter has information rents. One goal of the contract is to extract surplus from the recruiter and transfer it to the firm. This goal is generally in tension with the first-stage screening goal. Intuitively, trying to approximate the first-best suggestion region requires setting $\beta > 0$ and $\alpha > \bar{u}$, which means that the recruiter gets the outside option for the candidates on the boundary of the suggestion region but gets a strictly greater expected utility for all other suggested candidates. As we show in Section 6, although many specifications of the model allow the firm to extract all surplus, the first-best is still not achieved.

Commitment. When the firm can commit to a termination rule, γ is an additional degree of freedom that can be used to navigate the tensions outlined previously. However, when the firm cannot commit, $\gamma = \beta$ because the firm terminates low-productivity candidates in order to obtain the refund. In this case several equilibrium forces compete to determine β.

Recall that β relative to $\alpha - \bar{u}$ is a measure of the recruiter's incentive to screen candidates. With commitment, as β rises, the recruiter engages in more intense screening. Without commitment, increasing β has two additional impacts. First, it amplifies screening incentives because it increases not just the refund the recruiter has to pay but also the probability the recruiter has to pay the refund. Second, increasing β shifts the firm's payoff from individual candidates because the firm's ex post value function is $\max\{a, \beta\}$.

4 Parametric Analysis without Commitment

In this section, we make a parametric assumption under which there is a unique equilibrium contract without commitment, which we obtain in closed-form. We illustrate that the misalignment discussed in Section 3 occurs even though the firm designs the equilibrium contract to maximize total surplus.

Assumption 1 *The prior distribution of the candidate's productivity is* $a \sim Pareto(\bar{a}, k)$. *Conditional on productivity* a, *signals are drawn i.i.d. from a uniform distribution with minimum* 0 *and maximum* a. *Information type* i *has* τ_i *signals:* $x = (x_1, ...x_{\tau_i})$, *where* $\tau_1 > \tau_2 > ... > \tau_N$.

Conditional on observing signal x for a candidate with information type i, the recruiter's posterior belief about productivity is $Pareto(\max\{\bar{a}, \{x_t\}_{t=1}^{\tau_i}\}, \tau_i + k)$. [12] Under this information structure, the recruiter's posterior belief about a candidate depends only on the maximum of the observed signals and the number of signals. This information structure satisfies our nonparametric ordering assumption in Section 3.

The candidate's information type matters only in determining the number of signals observed about the candidate (τ_i). For this reason, we occasionally suppress the information type subscript i and consider an arbitrary information type with τ signals. The posterior distribution depends only on the maximum of observed signals and $\bar{a}$, an object we denote by $x_{max}^{\tau} := \max\{\bar{a}, \{x_t\}_{t=1}^{\tau}\}$. Whenever we observe a maximum signal of x_{max}^{τ}, we know for sure that the candidate's productivity is at least x_{max}^{τ}.

Given x_{max}^{τ}, the posterior mean productivity is equal to $\mathbb{E}[a|x_{max}^{\tau}] = \frac{\tau+k}{\tau+k-1} \cdot x_{max}^{\tau}$. For the same maximum signal x_{max}^{τ}, the posterior distribution keeps the same minimum (x_{max}^{τ}) regardless of τ, but the expected productivity decreases in τ via the shape parameter of the posterior distribution ($\tau + k$). Intuitively, if we observe more productivity signals with the same overall maximum, we become more pessimistic about the candidate's productivity.

4.1 First-Best Suggestion Region

Recall that $\mu^*(i)$ is the lowest-posterior-expectation candidate from group i that the firm suggests in the first-best. Under this information structure we have that

$$\mu^*(i) = \inf_m \{m | \mathbb{E}[\max\{a, 0\}|i, \mu = m] \geq c\} = c.$$

The firm suggests and hires only candidates it expects to be worth the hiring cost c. Thus the hiring bar in terms of expected productivity is the same across all information types, and the firm is endogenously risk-neutral. In terms of group-based statistical discrimination, the firm is also neutral, technically engaging in both types of statistical discrimination.

4.2 The Equilibrium Contract

Recall that $\mu_{\alpha,\beta,\gamma}(i)$ is the lowest-posterior-expectation candidate from group i whom the recruiter suggests under a refund contract with parameters (α, β, γ). Under this information

12. It is well-known in the statistics literature that a Pareto prior and uniform signals are a conjugate family, meaning beliefs remain Pareto after updating (Fink 1997). For completeness, we provide a proof in Appendix Section 8.1.

structure we have that

$$\mu_{\alpha,\beta,\gamma}(i) = \inf_m \left\{ m \middle| \Pr(a < \gamma | i, \mu = m) < \frac{\alpha - \bar{u}}{\beta} \right\} = \gamma \frac{\tau_i + k}{\tau_i + k - 1} \left(1 - \frac{\alpha - \bar{u}}{\beta} \right)^{\frac{1}{\tau_i + k}}.$$

Without commitment, the firm terminates all candidates with productivity below ($\gamma = \beta$), and we have that

$$\mu_{\alpha,\beta}(i) = \beta \frac{\tau_i + k}{\tau_i + k - 1} \left(1 - \frac{\alpha - \bar{u}}{\beta} \right)^{\frac{1}{\tau_i + k}}.$$

This expression shows that as the termination threshold (γ) rises, the suggestion threshold of the recruiter is scaled up for all information types. In general, the upfront fee must exceed to outside option ($\alpha > \bar{u}$), so as the refund β rises, the suggestion threshold rises.

Notice that $\frac{\tau_i + k}{\tau_i + k - 1}$ is increasing in i, while $\left(1 - \frac{\alpha - \bar{u}}{\beta} \right)^{\frac{1}{\tau_i + k}}$ is weakly decreasing in i. As a result, the direction of group-based statistical discrimination is ambiguous: $\mu_{\alpha,\beta}(i)$ can be either upward sloping or downward sloping in information type i. In the first-best, $\mu(i)$ is constant in i. This is the tension between first-stage screening, inefficient terminations, and surplus extraction outlined in Section 3.4. If the firm wants to approximate the shape of the first-best suggestion region, the firm must set $\alpha > \bar{u}$. However, this requires both leaving surplus for the recruiter and hiring candidates who are terminated with no realized benefit to the firm. There is a uniquely optimal way to navigate this trade-off.

Theorem 1 *The unique equilibrium contract is a refund contract with*

$$\beta^* = \frac{\mathbb{E}_\tau \left[\frac{\tau}{\tau + k} \right]}{\mathbb{E}_\tau \left[\frac{\tau}{\tau + k - 1} \right]} c, \qquad \alpha^* = \bar{u}.$$

A candidate is suggested if and only if the candidate's maximum signal exceeds β^*, that is, when the candidate's posterior expected productivity exceeds $\mu_{\alpha^*,\beta^*}(i) := \frac{\tau_i + k}{\tau_i + k - 1} \beta^*$.

Proof of Theorem 1.

- **Transfers are positive.** When specifying the contract space, we did not restrict the sign of the refund or the suggestion payment. We now show that in equilibrium, the transfer from the firm to the recruiter for suggestion is at least the outside option ($\alpha \geq \bar{u}$) and the transfer from the recruiter back to the firm is weakly positive ($\beta \geq 0$). Suppose for the sake of contradiction that in an optimal contract $\beta < 0$. Because the firm cannot commit to a termination rule, the firm terminates candidates only

when $a \geq \beta$. This means all candidates are not terminated if they were suggested by a recruiter. Anticipating this situation, the recruiter either suggests all candidates if $\alpha \geq \bar{u}$, yielding negative profit, or suggests no candidates if $\alpha \leq \bar{u}$, yielding at most 0 profit. The firm can obtain strictly positive profit by offering an alternative contract $\alpha' = \bar{u}, \beta' = c$, which leaves the recruiter with 0 surplus and induces the recruiter to suggest all candidates who have productivity above the hiring cost. This contradicts optimality.

Suppose for the sake of contradiction that in an optimal contract $\alpha < \bar{u}$. Because $\beta \geq 0$, even if the recruiter knows for sure the candidate will not be terminated, the recruiter will not suggest the candidate, because $\alpha - \beta \cdot 0 < \bar{u}$. Therefore, no candidates are suggested, yielding 0 profit for the firm. The firm can obtain strictly positive profit by offering the alternative contract $\alpha' = \bar{u}, \beta' = c$. This contradicts optimality.

- **Full surplus extraction.** We now argue that any equilibrium contract extracts all surplus from the recruiter, by setting $\alpha = \bar{u}$. In Appendix Section 8.2 we show that profit is strictly decreasing in α for all β. Therefore, $\alpha = \bar{u}$ in any profit-maximizing contract. When $\alpha = \bar{u}$,

$$\mu_{\alpha,\beta}(i) = \beta \frac{\tau_i + k}{\tau_i + k - 1}.$$

Note that the posterior mean maps to the maximum of τ signals in the following way: $\mu = \frac{\tau+k}{\tau+k-1} x^{\tau}_{max}$. Thus we have that the minimum suggested posterior expectation implies a minimum suggested maximum signal that is the same across all information types (x_{EQ}):

$$x_{EQ} = \frac{\tau_i + k - 1}{\tau_i + k} \mu_{\alpha,\beta}(i) = \beta \frac{\tau_i + k}{\tau_i + k - 1} \frac{\tau_i + k - 1}{\tau_i + k} = \beta.$$

Since only candidates with a maximum signal above x_{EQ} are suggested, this has the additional implication that only candidates who are retained for sure are suggested. Thus the recruiter's payoff is $\bar{u}$, all surplus is extracted, and the firm maximizes total surplus:

$$\mathbb{E}[\mathbb{I}\{x^{\tau}_{max} \geq \beta\}(\mathbb{I}\{a \geq \beta\}a + \mathbb{I}\{a \leq \beta\}\beta - c - \alpha)] = \mathbb{E}[\mathbb{I}\{x^{\tau}_{max} \geq x_{EQ}\}(a - c - \bar{u})]. \quad (1)$$

- **Optimal payments.** To complete the proof, we derive the unique refund payment (β^*) which maximizes (1). We do this by solving the first-order condition for β when $\alpha = \bar{u}$. We then show the second-order condition is satisfied at β^*. Appendix Section 8.2 contains these calculations.■

All three forces from Section 3.4 (first stage screening, second stage screening and surplus extraction) play a role in determining the unique equilibrium contract. By setting β^* to be the ratio of two expectations, the firm approximates the first-best suggestion region as best it can given the other two goals and the coarseness of the contract. By setting $\alpha = \bar{u}, \beta > 0$ the firm simultaneously extracts all surplus from the recruiter and eliminates inefficient terminations.

Strong screening incentives are key to full surplus extraction and elimination of inefficient terminations. The suggestion payment is exactly equal to the recruiter's outside option, and the refund is economically meaningful (close to the hiring cost). The recruiter suggests only candidates the recruiter is certain will not be terminated. Because screening incentives are strong, the condition for group statistical discrimination in favor of more information is satisfied:

$$\frac{\alpha^* - \bar{u}}{\beta^*} = 0 < \frac{1}{2}.$$

Corollary 1.1 *The firm engages in group statistical discrimination in favor of less information in the first-best, and the recruiter engages in group statistical discrimination in favor of more information under the unique equilibrium contract.*

Under the unique equilibrium contract, the directions of statistical discrimination are misaligned in the first-best and equilibrium. In the first-best, the firm suggests and hires all candidates with posterior expectations above the hiring cost. In equilibrium, the recruiter suggests and the firm hires candidates with posterior expectations that exceed group-specific hiring bars. The bar is higher for groups that are less well-understood by the recruiter.

Beyond refund contracts themselves, a pivotal ingredient in preventing achievement of the first-best is information heterogeneity across candidates. To see this, note that when there is no information heterogeneity (i.e., there is only one group with τ signals), we can drop the expectation operators on the numerator and denominator of β^*, yielding $\beta^* = \frac{\tau+k-1}{\tau+k}c$:

$$\mu_{\alpha^*,\beta^*}(i) = \beta^* \frac{\tau + k}{\tau + k - 1} \frac{\tau + k - 1}{\tau + k} c = c.$$

Thus when the pool of candidates is homogeneous in terms of information quality, the recruiter suggests only candidates with posterior expectations above the hiring cost, and the first-best is achieved.[13]

13. There is full surplus extraction, and the first-best outcome is generally achieved in equilibrium when there is only one information type.

4.3 Hiring Outcomes by Group

Because the Pareto-uniform information structure without commitment yields a unique closed-form equilibrium contract, we can make stronger statements about how delegation impacts candidate groups differently. We present and discuss these results in this section, after defining the following cutoff.

Definition 3 *The cutoff number of signals τ^* is such that*

$$\frac{\tau^* + k - 1}{\tau^* + k} c = \beta^*.$$

A few observations follow immediately: (1) any information type with τ_i equal to τ^* will be hired at the same rates in the first-best and equilibrium; (2) τ^* is always unique and well-defined; (3) τ^* can be non-integer, but the analysis can be generalized to $\tau_i \in \mathbb{R}_+$, via a completion of the analysis for integer $\tau_i \in \mathbb{Z}_+$; (4) τ^* is strictly between τ_N and τ_1.

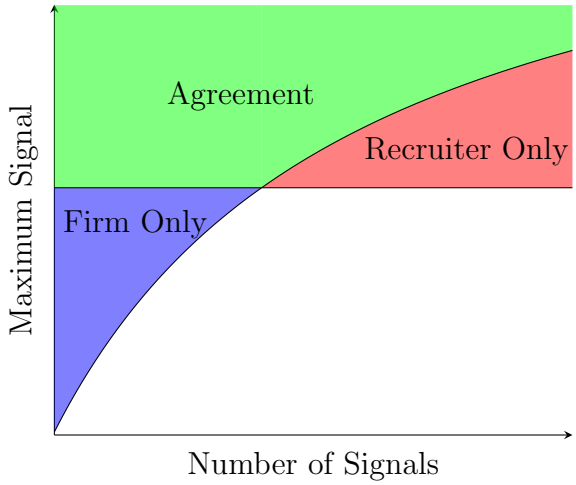
The equilibrium contract from Theorem 1 and the first-best hiring threshold generate misalignment between the two suggestion regions. Further, delegation systematically improves the hiring outcomes of groups of candidates who are better understood by recruiters.

Theorem 2 *All information types with more signals than τ^*, i.e., $i < \min\{i' | \tau_{i'} \leq \tau^*\}$, are hired with higher probability in equilibrium than in the first-best, while all information types with fewer signals than τ^*, i.e., $i > \max\{i' | \tau_{i'} \geq \tau^*\}$, are hired with lower probability than in the first-best.*

The complete proof is provided in Appendix Section 8.4. Under the equilibrium contract, there is a common threshold for hired candidates in terms of the maximum observed signal: the refund payment (β^*). This corresponds to the horizontal line in Figure 2. We show via induction that this threshold is both strictly below the first-best threshold for the information type with the most signals and strictly above the first-best threshold for the information type with the least signals. This implies that the blue and red regions in Figure 2 are not empty.

There is a positive share of candidates who are not hired in equilibrium but that the firm would like to hire because their maximum signal is high enough given the small number of signals observed. These candidates are depicted as the "Firm Only" (blue) region in Figure 2. There is also a positive share of candidates who are hired in equilibrium that the firm would prefer not to hire because their maximum signal is not high enough given the large number of signals observed. These candidates are depicted as the "Recruiter Only" (red) region in Figure 2.

Figure 2: Types of Candidates Hired

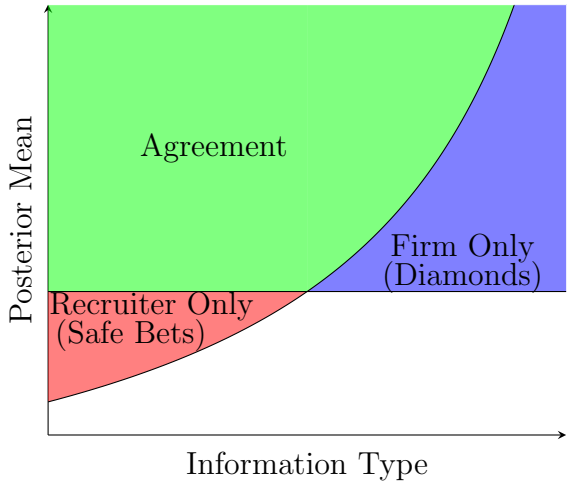


Note: The space of candidates is depicted in terms of their information type (number of signals) and expected productivity. The green region includes all candidates hired in both the first-best and equilibrium. The red region includes candidates hired only in equilibrium. The blue region includes candidates hired only in the first-best.

In the first-best benchmark only candidates with posterior expected productivity greater than the hiring cost (c) are hired. In equilibrium, the trade-offs induced by delegation are such that the firm optimally sets the suggestion bar to a common threshold in terms of signals. Because different signal realizations imply different expected productivity across information types, the recruiter suggests some candidates with posterior expected productivity less than c and does not suggest some candidates with expected productivity above c. First-stage screening is distorted in equilibrium relative to the first-best.

To better understand this result, in Figure 3 we depict the set of suggested and hired candidates in terms of their information types and posterior means, (i, μ). The key misalignment is that the equilibrium suggestion region excludes candidates who have high posterior expected productivity but a high residual uncertainty in less informative groups (so-called diamonds in the rough, depicted in blue in Figure 3), and includes candidates who have low posterior expected productivity and a low residual uncertainty from more informative groups (so-called safe bets, depicted in red in Figure 3). The optimal contract always generates a positive share of the two types of candidates.

Figure 3: Candidate Types Hired in Equilibrium vs. First-Best



Note: Candidates are depicted in terms of their posterior expectation and information type. Diamonds in the rough are candidates that are poorly understood (high information type) but are expected to have high productivity. Safe bets are candidates that are well-understood (many signals) but are expected to have low productivity. The green region includes all candidates hired in both the first-best and equilibrium. The red region includes candidates hired only in equilibrium. The blue region includes candidates hired only in the first-best.

5 Comparative Statics

In the baseline model, underlying productivity ($\bar{a}$) is distributed similarly across information types. In this section we relax this assumption and allow candidate types to vary in the scale parameter of their prior productivity ($\bar{a}_i$). We also ask how improvements in information, group size, and productivity of one information type impact the other types. We show that in the first-best, hiring outcomes of information types are independent, while in equilibrium delegation via refund contracts causes spillovers.

We begin by considering the first-best. Following the same logic as in the baseline model, in the first-best, the firm hires a candidate if $\mu^*(i) = c$.

Theorem 1' *The unique equilibrium contract is a refund contract with*

$$\beta^* = \frac{\mathbb{E}\left[\frac{\bar{a}_i^k \tau_i}{\tau_i + k}\right]}{\mathbb{E}\left[\frac{\bar{a}_i^k \tau_i}{\tau_i + k - 1}\right]} c, \qquad \alpha^* = \bar{u}.$$

The proof of Theorem 1' requires only minor additions to the proof of Theorem 1 and is

given in Appendix Section 8.3. A key observation is that heterogeneity in prior productivity essentially re-weights information types within the expectations that determine the optimal refund. Notice that $\frac{\alpha^*-\bar{u}}{\beta^*} = 0 < \frac{1}{2}$ and thus we have group statistical discrimination in favor of more information as in the baseline model. We now analyze how an increase in information for one information type affects the other types in the population.

Proposition 3 *Suppose more signals become available for information type i with $\tau_i > \tau^*$, i.e., τ_i increases. Then, β^* increases and candidates from any type that is not i are hired with lower unconditional probability (for all other information types together, each type separately, or any given candidate of any ability and of any type that is not i).*

The proof is provided in Appendix Section 8.5.1. In equilibrium, additional information for one information type spills over onto other information types because it raises the refund payment. In the first-best, each type has its own threshold signal and information improvements for one type do not impact the other types. Consider the case when there are only two information types, one with many signals and one with few. If we add a signal to the information type that already has many signals, we increase the threshold we require of all types, further widening the difference in hiring probabilities between the first-best and equilibrium for the low type. This happens because of two forces acting in the same direction: (1) more signals increase the individual first-best threshold for this type, increasing the "average" equilibrium threshold; (2) by improving information for the type, we make correctly screening the type more important, which means that we should put "more weight" on this individual type first-best threshold, which is above the "average" equilibrium threshold. In this sense, we increase delegation-based statistical discrimination.

Second, we analyze how first-order stochastic improvements of ability in one information type affect the hiring probabilities for all other information types. The proofs for the remaining two results are provided in Appendix Section 8.5.2.

Proposition 4 *Suppose the prior productivity of information type i improves, i.e., $\bar{a}_i$ increases. If τ_i is greater (less) than τ^*, then β^* increases (decreases) and candidates from any other type (not i) are hired with lower (higher) unconditional probabilities (for all other information types together, each type separately, or any given candidate of any productivity and of any type that is not i).*

This comparative static highlights that the equilibrium contract is always attempting to balance the inclusion of safe bets and the exclusion of diamonds in the rough. In any equilibrium, diamonds in the rough come exclusively from information types below the cutoff τ^*, while safe bets come exclusively from information types above the cutoff. A productivity

improvement for an information type below the cutoff makes excluding diamonds in the rough more expensive, and the firm responds by reducing β in order to reduce screening. A productivity improvement for an information type above the cutoff makes including safe bets more expensive, and the firm responds by increasing β in order to improve screening.

Finally, we analyze how a change in the size of an information type impacts the hiring probabilities of the other information types.

Proposition 5 *Assume that information type i becomes larger relative to the other information types, i.e., p_i increases and all other p_j proportionally decrease. If τ_i is greater (less) than τ^*, then β^* increases (decreases) and candidates from any other type (not i) are hired with lower (higher) unconditional probabilities (for all other types together, each type separately, or any given candidate of any productivity and of any type that is not i).*

The similarity between Propositions 4 and 5 illustrates that productivity improvements have similar effects to changing the relative sizes of information types. The intuition for this result is similar in spirit to Proposition 4. Increasing the size of information types below the cutoff means there are more diamonds in the rough, so the firm reduces screening. Increasing the size of information types above the cutoff means that there are more safe bets, so the firm improves screening.

6 Robustness and Extensions

In this section, we show that the qualitative results derived under the parametric information structure without commitment are robust. We do this by considering several variations of the model. Without commitment, if the contract space is enriched to include a retainer payment which can be used to extract all surplus, the first-best is still not achieved. When the recruiter has a choice over an intensive margin of screening (modeled using costly sequential search), misalignment and different types of statistical discrimination still arise. With commitment, full extraction generally occurs, but misalignment persists due to the fundamental tension between first- and second-stage screening.

6.1 Three-Part Contracts without Commitment

In this subsection we consider the case without commitment ($\gamma = \beta$), and we show the need to use (α, β) to extract surplus is not pivotal in preventing the achievement of the first-best. To do this, we enrich the space of contracts to include an additional transfer which is paid out prior to the recruiter observing the information type and productivity signals of candidates.

This additional transfer functions as a form of retainer, and it can be used directly by the firm to extract all surplus from the recruiter, freeing up other parts of the contract (α, β) to play other roles. We continue to assume the Pareto-uniform information structure for this result.

Proposition 6 *Suppose the firm designs three-part contracts, with an additional transfer before the recruiter sees the productivity signals. Then the following are true:*

1. *Profit is weakly higher than in the baseline equilibrium.*
2. *The first-best profit and set of hired and suggested candidates are not achieved.*

Proof. Recall that the recruiter cannot commit to a suggestion decision rule, because the signals are not contractible. In the model we consider, $\bar{u}$ is the outside option for rejecting the contract and also the outside option for not suggesting any candidate after the contract is accepted. So, if the recruiter accepts the contract but does not suggest a candidate, the recruiter's payoff is the upfront transfer plus $\bar{u}$. The firm can replicate the profit from the baseline model by setting the additional transfer to be 0 and setting the other contract payments to be the same. Therefore, the firm's profit is weakly higher.

The additional transfer occurs before the signals are realized, so it impacts only whether the recruiter accepts the contract. In any equilibrium, the firm sets this transfer such that the recruiter's expected utility is equal to the recruiter's outside option $\bar{u}$. This implies the firm maximizes total surplus in order to maximize profit. Since total surplus depends only on the set of candidates suggested and terminated, achieving the first-best requires achieving the same set of suggested and terminated candidates.

Suppose for the sake of contradiction that the first-best hiring and suggestion regions are achieved. In the first-best, all candidates who are suggested are not terminated. To achieve this scenario in equilibrium, the recruiter must not suggest anyone who is later terminated. This requires $\beta > 0$ and the upfront payment equal to the outside option $\alpha = \bar{u}$. However, this implies that the suggestion threshold in terms of signals $x^{\tau}_{EQ}(\alpha, \beta) = \beta$, which is the same for all information types. This is a contradiction, because in the first-best the suggestion thresholds are different for different types: $x^{\tau}_{FB} = \frac{\tau+k-1}{\tau+k}c$.■

The proposition demonstrates that the need to extract rent is not pivotal in terms of preventing achievement of the first-best. The proof of the proposition further confirms that the deeper tension is between achieving the first-best termination level and achieving the first-best suggestion region.

6.2 Multiple Candidates and Delegated Search

In our baseline model, the recruiter either suggests the candidate or not. This situation, which can be viewed as screening along an extensive margin, captures many realistic scenarios where the recruiter has a candidate in hand who they can choose to suggest for an opportunity. In other situations, the recruiter may have the ability to search more or less hard for a candidate that will satisfy the firm's hiring requirements. This situation can be viewed as screening along an intensive margin.

Consider a model where the recruiter can pay a search cost to sample additional candidates. After sampling a candidate, the recruiter views the candidate's information type and signals. In such a situation, the refund contract determines not just the set of suggested candidates but also the amount of costly effort the recruiter spends searching for candidates. How does this impact misalignment and statistical discrimination?

In such a sequential search model with normal signals and two information types, the spirit of our results continues to hold. The candidate from the more uncertain group is hired more often in the first-best and equilibrium. Intuitively, adding an intensive margin introduces moral hazard issues. The contract must have strong incentives not just so that the right candidates are chosen but also so that enough search effort is exerted by the recruiter. In general, the firm will choose to encourage less search in this setting because the selection of candidates is distorted.

6.3 Nonparametric Full Surplus Extraction with Commitment

In this subsection, we show the equilibrium contract with commitment extracts all surplus from the recruiter for any information structure which satisfies our nonparametric ordering assumption. We also discuss how despite full surplus extraction, the tension between first- and second-stage screening generally persists and prevents achievement of the first-best. This illustrates that the full extraction and failure to achieve the first-best that we observe in our parametric analysis without commitment is a robust feature of equilibrium contracts.

Proposition 7 *The optimal contract with commitment features full extraction.* [14]

Proof

- **Transfers are positive.** We prove that in equilibrium, the transfer from the firm to the recruiter for suggestion of a candidate is at least the outside option ($\alpha \geq \bar{u}$), and the

14. There is a sequence of contracts that converges to the firm's supremum profit and features full extraction in the limit.

transfer from the recruiter back to the firm is weakly positive ($\beta \geq 0$) by contradiction. Suppose that in an optimal contract $\beta < 0$. Then the recruiter will suggest a candidate with termination probability above threshold $p^* = \frac{\alpha - \bar{u}}{\beta}$, or $\mu < \mu_i$, which yields a negative selection among suggested candidates and contradicts optimality.

If β=0, then the recruiter either suggests any candidate if $\alpha \geq \bar{u}$ or suggests no candidate otherwise. Either situation is suboptimal by assumption. If $\alpha < \bar{u}$ and $\beta \geq 0$, the recruiter will never suggest the candidate, because $\alpha - \beta \Pr(a < \gamma | i, \mu) < \bar{u}$. Thus, the contract yields 0 profit for the firm, which again contradicts optimality.

- **Refund is equal to zero and the firm extracts the total surplus.** Assume any contract with $\beta > 0$, $\alpha > \bar{u}$, and some γ. The recruiter suggests all candidates with the equilibrium termination probability $\Pr(a < \gamma | i, \mu) \leq p^* = \frac{\alpha - \bar{u}}{\beta}$ and gets strictly positive expected utility for those candidates with termination probability strictly below p^*. We propose increasing the firm's profit by a sequence of contracts with the same suggestion and termination rules (and thus the surplus) that converges to full extraction of the surplus from the recruiter:

$$(\beta_m,\ \alpha_m = \bar{u} + \beta_m \cdot \frac{\alpha - \bar{u}}{\beta},\ \gamma_m = \gamma)_{m=1}^{\infty}$$

and $\lim_{m \to \infty} \beta_m = 0$ ($\beta_m > 0\ \forall m$). Each of these contracts has $p_m^* = p^*$ and thus leads to the same suggesting decisions. Moreover, γ is fixed across all contracts, which means that the sets of hired and the set of retained candidates do not change. Finally, the recruiter's overall expected utility decreases and converges to $\bar{u}$ (as it is bounded between $\bar{u}$ and α), which, given the fixed surplus, implies that the firm's expected payoff increases. ■

The proposition confirms that under commitment, all equilibria feature full surplus extraction. The proof holds economic insight: The firm can vary β, α in such a way that it fully extracts all surplus while maintaining the same set of suggested candidates.

We can interpret this result in light of the three goals of the firm outlined in Section 3.4: first-stage screening, second-stage screening, and surplus extraction. Proposition 7 implies that with commitment the firm can perfectly extract all surplus and still be left with two degrees of freedom to accomplish the other two goals. However, the remaining two goals are still in conflict. Replicating the first-best suggestion region generally requires inefficient terminations compared to the first-best. Thus, the first-best will not generally be achieved even with commitment.

7 Conclusion

This paper asks how refund contracts, which are used to delegate recruitment to recruiters, shape the types of candidates ultimately hired by the firm. We show that relative to a first-best benchmark where a firm recruits directly, delegation induces endogenous risk aversion. Refund contracts make the recruiter less likely to suggest and the firm less likely to hire candidates that are less well-understood by the recruiter.

We call this statistical discrimination in favor of more information. We show that it occurs under many refund contracts and many variants of the model. In the first-best, the firm generally statistically discriminates in favor of less information, because it can terminate poor performers. Despite this result, we show that in one tractable case, the unique equilibrium contract features misalignment. Thus, the firm is typically unable to design a contract which avoids misalignment of statistical discrimination.

We show that the presence of information differences across candidates is a main driver of misalignment. When candidates are all understood equally well, the first-best can be achieved using refund contracts. This could give rise to a vicious cycle, the nature of which we sketch here but leave the study of for future work. Suppose candidates that are not hired get discouraged and exit an industry. Suppose also that firms can recruit directly at an opportunity cost or hire a recruiter. Because candidates about whom recruiters are less informed are less likely to be hired, they will exit at higher rates. In subsequent periods, the pool of candidates will be more homogeneous, and delegation will become attractive for a greater share of firms. This would then encourage more firms to delegate, thus continuing the cycle.

We see two veins for future work. First, the theoretical impact of refund contracts should be tested empirically using a two-sided audit study, where recruiters are hired to fill a position, and the experiment varies both the types of candidates the recruiters have access to and the type of compensation scheme. Second, how candidates behave in the presence of widespread delegation should be explored. Job candidates often make costly choices about which signals to acquire, and these choices are shaped by the labor market returns to these signals. We have shown in this paper that delegated recruitment changes the return to additional signals. It remains unclear if job candidates strategically change how they build their resumes and careers as recruiters and headhunters become more common.

References

Arslanian, Henri. 2016. "Talent and Hiring." *Entrepreneurship in Finance: Successfully Launching and Managing a Hedge Fund in Asia,* 183–193.

Association, American Staffing. 2014. September. https://americanstaffing.net/wp-content/uploads/2014/09/NAPS_ASA_recruiting_agrmnt.doc.

Black, Ines, Sharique Hasan, and Rembrand Koning. 2022. "Hunting for talent: Firm-driven labor market search in the United States." *Available at SSRN.*

Carroll, Gabriel. 2017. "Robustness and separation in multidimensional screening." *Econometrica* 85 (2): 453–488.

Chade, Hector, and Natalia Kovrijnykh. 2016. "Delegated information acquisition with moral hazard." *Journal of Economic Theory* 162:55–92.

Che, Yeon-Koo, Wouter Dessein, and Navin Kartik. 2013. "Pandering to persuade." *American Economic Review* 103 (1): 47–79.

Cowgill, Bo, and Patryk Perkowski. 2020. "Delegation in Hiring: Evidence from a Two-Sided Audit." *Columbia Business School Research Paper,* no. 898.

Deutsch, Matthew. 2019. *Most popular type of Recruitment Guarantee: Agency Guarantee period,* March. https://www.topechelon.com/blog/recruiter-guarantees-once-i-am-paid-i-do-not-refund-any-money/.

Dingman, H Bruce. 1993. "Contingency recruiting in the hospitality industry." *Cornell Hospitality Quarterly* 34 (6): 30.

EEOC. 2023. "Prohibited Employment Policies/Practices," August. https://www.eeoc.gov/prohibited-employment-policiespractices.

Fink, Daniel. 1997. "A compendium of conjugate priors." *See http://www. people. cornell. edu/pages/df36/CONJINTRnew% 20TEX. pdf* 46.

Finlay, William, and James E Coverdill. 2007. *Headhunters: Matchmaking in the labor market.* Cornell University Press.

Florea, Nicoleta Valentina. 2014. "Using recruitment agencies to obtain the best candidates." *Land Forces Academy Review* 19 (1): 80.

Frankel, Alexander. 2014. "Aligned delegation." *American Economic Review* 104 (1): 66–83.

———. 2021. "Selecting Applicants." *Econometrica* 89 (2): 615–645.

Gale, Douglas, and Martin Hellwig. 1985. "Incentive-compatible debt contracts: The one-period problem." *The Review of Economic Studies* 52 (4): 647–663.

Giuliano, Laura, David I Levine, and Jonathan Leonard. 2009. "Manager race and the race of new hires." *Journal of Labor Economics* 27 (4): 589–631.

Inderst, Roman, and Marco Ottaviani. 2012. "Competition through commissions and kickbacks." *American Economic Review* 102 (2): 780–809.

Kundu, Tapas, and Tore Nilssen. 2020. "Delegation of regulation." *The Journal of Industrial Economics* 68 (3): 445–482.

Levitt, Steven D, and Christopher M Snyder. 1997. "Is no news bad news? Information transmission and the role of" early warning" in the principal-agent model." *The RAND Journal of Economics,* 641–661.

Perry, David E, and Mark J Haluska. 2017. *Executive Recruiting For Dummies.* John Wiley & Sons.

Press, Entrepreneur. 2007. *Start Your Own Executive Recruiting Service: Your Step-By-Step Guide to Success.* Entrepreneur Press.

Steinberg, Sheryl L, and David S Machlowitz. 1989. "How to Work with a Headhunter." *Legal Econ.* 15:43.

Szalay, Dezsö. 2005. "The economics of clear advice and extreme options." *The Review of Economic Studies* 72 (4): 1173–1198.

———. 2009. "Contracts with endogenous information." *Games and Economic Behavior* 65 (2): 586–625.

Yang, Frank. 2021. "Costly multidimensional screening." *arXiv preprint arXiv:2109.00487.*

8 Appendix

8.1 Bayesian Updating Under Pareto-Uniform Information Structure

To derive the posterior distribution, note that $x|a \sim U[0,a]$. Therefore the joint probability density function of τ signals given a is:

$$\prod_{t=1}^{\tau} f_{x|a}(x_t) = \frac{1}{a^\tau} \prod_{t=1}^{\tau} \mathbb{I}\{x_t \leq a\}$$

Notice that $\prod_{t=1}^{\tau} \mathbb{I}\{x_t \leq a\} = 1$ if and only if $x^\tau_{max} \leq a$. Therefore we can re-write as:

$$\prod_{t=1}^{\tau} f_{x|a}(x_t) = \frac{1}{a^\tau} \mathbb{I}\{x^\tau_{max} \leq a\}$$

The joint distribution of $a, \{x_t\}_{t=1}^{\tau}$ is then the product of this distribution and the prior distribution of a:

$$f(a, \{x_t\}_{t=1}^{\tau}) = \frac{1}{a^\tau} \mathbb{I}\{x^\tau_{max} \leq a\} \mathbb{I}\{a \geq \bar{a}\} \frac{k\bar{a}^k}{a^{k+1}} \tag{2}$$

$$= \frac{k\bar{a}^k}{a^{\tau+k+1}} \mathbb{I}\{\max\{x^\tau_{max}, \bar{a}\} \leq a\} \tag{3}$$

Notice that this depends only on the maximum signal. Therefore x^τ_{max} is a sufficient statistic for $\{x_t\}_{t=1}^{\tau}$ for a. Note that the conditional PDF will be proportional to the joint PDF for any fixed x^τ_{max}. We can then solve for the multiplicative constant A that makes the conditional PDF integrate to 1 over the support:

$$f(a|\{x_t\}_{t=1}^{\tau}) = \int_{\max\{x^\tau_{max}, \bar{a}\}}^{\infty} \frac{A}{a^{\tau+k+1}} da \leftrightarrow \frac{A}{(\tau+k)\max\{x^\tau_{max}, \bar{a}\}^{\tau+k}} = 1$$

$$A = (\tau+k)\max\{x^\tau_{max}, \bar{a}\}^{\tau+k} \leftrightarrow f(a|\{x_t\}_{t=1}^{\tau}) = \mathbb{I}\{\max\{x^\tau_{max}, \bar{a}\} \leq a\} \frac{(\tau+k)\max\{x^\tau_{max}, \bar{a}\}^{\tau+k}}{a^{\tau+k+1}}$$

Thus, $a|\{x_t\}_{t=1}^{\tau} \sim Pareto(\max\{x^\tau_{max}, \bar{a}\}, \tau+k)$.

8.2 Algebra for the Proof of Theorem 1

This section proves that under refund contracts, firm profit is strictly decreasing in α. It then shows that after $\alpha = \bar{u}$, β^* given in the theorem solves the first-order condition. It concludes by showing that β^* satisfies the second-order condition. For historical reasons and to reduce

algebra, derivations are performed by making the change of variables: $\tilde{\alpha} = \beta - \alpha + \bar{u}, \tilde{\beta} = \beta$. $\tilde{\beta}$ is a bonus, and $\tilde{\alpha}$ is an upfront payment incorporating the outside option. The recruiter suggests a candidate if $\tilde{\beta} Pr(a \geq \beta | x^{\tau}_{max}) - \tilde{\alpha} \geq 0$, where note $\tilde{\alpha}$ is taken from the recruiter. The restriction that α be weakly greater than the outside option translates to $\beta \geq \tilde{\alpha}$. We now suppress the tilde notation. The firm chooses a contract to maximize profit:

$$\max_{\beta \geq 0, \alpha \leq \beta,} \mathbb{E}\left[\mathbb{I}\{x^{\tau}_{max} \geq x^{\tau}_{EQ}(\alpha, \beta)\}\left(\mathbb{I}\{a \geq \beta\}(a - \beta) - c + \alpha\right)\right]$$

This can be broken into three parts. First, the expected productivity benefit of hiring a candidate:

$$\mathbb{E}[\mathbb{I}\{x^{\tau}_{max} \geq x^{\tau}_{EQ}(\alpha, \beta)\}\mathbb{I}\{a \geq \beta\}a] \tag{4}$$

Second, the expected suggestion costs:

$$\mathbb{E}[\mathbb{I}\{x^{\tau}_{max} \geq x^{\tau}_{EQ}(\alpha, \beta)\}(\alpha - c)] \tag{5}$$

Third, the expected bonus cost:

$$-\mathbb{E}[\mathbb{I}\{x^{\tau}_{max} \geq x^{\tau}_{EQ}(\alpha, \beta)\}\mathbb{I}\{a \geq \beta\}\beta] \tag{6}$$

The productivity benefit (4) can be written as:

$$E[\mathbb{I}\{a > \beta\}\mathbb{I}\{x^{max}_{\tau} \geq \beta\left(\frac{\alpha}{\beta}\right)^{\frac{1}{\tau+k}}\}a] = \frac{\tau}{\tau + k - 1}\frac{k\bar{a}^k}{k-1}\beta^{1-k} + \frac{k\bar{a}^k}{\tau + k - 1}\beta^{1-\tau-k}\left[\beta^{\tau} - \left(\beta\left(\frac{\alpha}{\beta}\right)^{\frac{1}{\tau+k}}\right)^{\tau}\right] \tag{7}$$

$$= \bar{a}^k \frac{k}{\tau + k - 1}\beta^{1-k}\left[\frac{\tau}{k-1} + 1 - \left(\frac{\alpha}{\beta}\right)^{\frac{\tau}{\tau+k}}\right] \tag{8}$$

$$= \beta\left(\frac{\bar{a}}{\beta}\right)^k \frac{k}{\tau + k - 1}\left[\frac{\tau}{k-1} + 1 - \left(\frac{\alpha}{\beta}\right)^{\frac{\tau}{\tau+k}}\right] \tag{9}$$

The expected suggestions costs (5) can be written as:

$$(-c + \alpha)Pr(x^{max}_{\tau} \geq \beta\left(\frac{\alpha}{\beta}\right)^{\frac{1}{\tau+k}}) = \frac{\tau}{\tau + k}\left(\frac{\bar{a}}{\alpha^{\frac{1}{\tau+k}}}\right)^k \beta^{\frac{k}{\tau+k}-k}(-c + \alpha) \tag{10}$$

$$= \frac{\tau}{\tau + k}\left(\frac{\bar{a}}{\beta}\right)^k \left(\frac{\beta}{\alpha}\right)^{\frac{k}{\tau+k}}(-c + \alpha) \tag{11}$$

$$= \left(\frac{\alpha}{\beta}\right)^{\frac{\tau}{\tau+k}} \beta\left(\frac{\bar{a}}{\beta}\right)^k \frac{\tau}{\tau + k} - c\frac{\tau}{\tau + k}\left(\frac{\bar{a}}{\beta}\right)^k \left(\frac{\beta}{\alpha}\right)^{\frac{k}{\tau+k}} \tag{12}$$

The expected bonus cost (6) can be written as:

$$-\beta Pr(a \geq \beta, x_\tau^{max} \geq x_{EQ}^\tau(\alpha,\beta)) = -\beta Pr(x_{max}^\tau \geq \beta) - \beta Pr(x_{EQ}^\tau(\alpha,\beta) \leq x_\tau^{max} \leq \beta, a \geq \beta) \tag{13}$$

$$= -\beta \frac{\tau}{\tau+k}\left(\frac{\bar{a}}{\beta}\right)^k - \beta \int_{\beta(\frac{\alpha}{\beta})^{\frac{1}{\tau+k}}}^{\beta} \left(\frac{x}{\beta}\right)^{\tau+k} \frac{k\tau\bar{a}^k}{(\tau+k)x^{k+1}} dx \tag{14}$$

$$= -\beta \frac{\tau}{\tau+k}\left(\frac{\bar{a}}{\beta}\right)^k - \frac{\tau\beta k\bar{a}^k}{(\tau+k)\beta^{\tau+k}} \int_{\beta(\frac{\alpha}{\beta})^{\frac{1}{\tau+k}}}^{\beta} x^{\tau-1} dx \tag{15}$$

$$= -\beta \frac{\tau}{\tau+k}\left(\frac{\bar{a}}{\beta}\right)^k - \frac{\beta k}{\tau+k}\left(\frac{\bar{a}}{\beta}\right)^k \left(1 - \left(\frac{\alpha}{\beta}\right)^{\frac{\tau}{\tau+k}}\right) \tag{16}$$

$$= \beta\left(\frac{\bar{a}}{\beta}\right)^k \left[-1 + \frac{k}{\tau+k}\left(\frac{\alpha}{\beta}\right)^{\frac{\tau}{\tau+k}}\right] \tag{17}$$

Profit is the sum of (9), (12), and (17):

$$\beta\left(\frac{\bar{a}}{\beta}\right)^k \left[\left(\frac{\alpha}{\beta}\right)^{\frac{\tau}{\tau+k}} \frac{\tau-1}{\tau+k-1} + \frac{1}{k-1} - \frac{c}{\beta}\frac{\tau}{\tau+k}\left(\frac{\beta}{\alpha}\right)^{\frac{k}{\tau+k}}\right]$$

Note that this expression is strictly increasing in α, and it is strictly increasing in α for all valid values of τ, therefore total profit is increasing in α. Thus in the optimal contract the upfront payment is set to its maximal possible value: $\alpha = \beta$. The expression now becomes only a function of β:

$$\beta\left(\frac{\bar{a}}{\beta}\right)^k \left[\frac{\tau-1}{\tau+k-1} + \frac{1}{k-1} - \frac{c}{\beta}\frac{\tau}{\tau+k}\right]$$

Simplifying:

$$\pi_i = \left(\frac{\bar{a}}{\beta}\right)^k \left[\frac{\tau k}{(k-1)(\tau+k-1)}\beta - \frac{c\tau}{\tau+k}\right] \tag{18}$$

The first order condition for profit from information type i is:

$$\frac{\partial \pi_i}{\partial \beta} = \frac{-k}{\beta}\left(\frac{\bar{a}}{\beta}\right)^k\left[\frac{\tau k}{(k-1)(\tau+k-1)}\beta - \frac{c\tau}{\tau+k}\right] + \left(\frac{\bar{a}}{\beta}\right)^k \frac{\tau k}{(k-1)(\tau+k-1)} \tag{19}$$

$$= k\left(\frac{\bar{a}}{\beta}\right)^k\left\{\frac{c}{\beta}\frac{\tau}{\tau+k} + \frac{\tau}{(k-1)(\tau+k-1)} - \frac{\tau k}{(k-1)(\tau+k-1)}\right\} \tag{20}$$

$$= k\left(\frac{\bar{a}}{\beta}\right)^k\left\{\frac{c}{\beta}\frac{\tau}{\tau+k} - \frac{\tau(k-1)}{(k-1)(\tau+k-1)}\right\} \tag{21}$$

$$= k\left(\frac{\bar{a}}{\beta}\right)^k\left\{\frac{c}{\beta}\frac{\tau}{\tau+k} - \frac{\tau}{\tau+k-1}\right\} \tag{22}$$

$$\tag{23}$$

$$\frac{\partial \pi_i}{\partial \beta} = k\left(\frac{\bar{a}}{\beta}\right)^k\left\{\frac{c}{\beta}\frac{\tau}{\tau+k} - \frac{\tau}{\tau+k-1}\right\} \tag{24}$$

Taking the weighted average of the first order condition for each information type gives the total profit first-order condition:

$$\sum_i p_i k\left(\frac{\bar{a}}{\beta}\right)^k\left[-\frac{\tau_i}{\tau_i+k-1} + \frac{c}{\beta}\frac{\tau_i}{\tau_i+k}\right] = k\left(\frac{\bar{a}}{\beta}\right)^k\left[-\sum_i p_i\frac{\tau_i}{\tau_i+k-1} + \frac{c}{\beta}\sum_i p_i\frac{\tau_i}{\tau_i+k}\right]$$

$$= k\left(\frac{\bar{a}}{\beta}\right)^k\left(-\mathbb{E}\left[\frac{\tau_i}{\tau_i+k-1}\right] + \frac{c}{\beta}\mathbb{E}\left[\frac{\tau_i}{\tau_i+k}\right]\right) = 0$$

Which yields the optimal value of β:

$$\beta = \frac{\mathbb{E}\left[\frac{\tau_i}{\tau_i+k}\right]}{\mathbb{E}\left[\frac{\tau_i}{\tau_i+k-1}\right]}c$$

Returning to our original change of variables, $\tilde{\beta} = \beta$ and $\tilde{\alpha} = \tilde{\beta}$ thus $\alpha = \bar{u}$. We show the second-order condition is satisfied at β^* for the case with prior productivity heterogeneity (which nests the baseline model) in the next section.

8.3 Optimal Contract with Prior Heterogeneity

Suppose additionally that information types also differ in prior productivity. That is they differ in their Pareto shift parameter $\bar{a}_i$ in addition to the number of signals. We can still use the same argument as above to say that $\alpha = \beta$ because the argument is true for each

information type individually regardless of the values of $\bar{a}, k$. We can adapt the first-order condition (24) to incorporate heterogeneity:

$$\sum_i p_i k \left(\frac{\bar{a}_i}{\beta}\right)^k \left\{ \frac{c}{\beta} \frac{\tau_i}{\tau_i + k} - \frac{\tau_i}{\tau_i + k - 1} \right\} = 0 \tag{25}$$

$$\sum_i p_i \left[\frac{\bar{a}_i^k \tau_i}{\tau_i + k - 1} \beta - \frac{\bar{a}_i^k c \tau_i}{\tau_i + k} \right] = 0 \tag{26}$$

$$\beta \sum_i p_i \frac{\bar{a}_i^k \tau_i}{\tau_i + k - 1} = \sum_i p_i \frac{\bar{a}_i^k c \tau_i}{\tau_i + k} \tag{27}$$

$$\beta^* = \frac{\sum_i p_i \frac{\bar{a}_i^k c \tau_i}{\tau_i + k}}{\sum_i p_i \frac{\bar{a}_i^k \tau_i}{\tau_i + k - 1}} \tag{28}$$

$$\beta^* = \frac{\mathbb{E}\left[\frac{\bar{a}_i^k \tau_i}{\tau_i + k}\right]}{\mathbb{E}\left[\frac{\bar{a}_i^k \tau_i}{\tau_i + k - 1}\right]} c \tag{29}$$

We can also check the second-order condition, to make sure that the solution is a local maximum. It is then also the global maximum, since the function is everywhere differentiable, and the solution to the first-order condition is unique.

$$\frac{\partial^2 \pi}{\partial \beta^2} = \sum_i \frac{p_i k \tau_i}{\beta} \cdot \left(\frac{\bar{a}_i}{\beta}\right)^k \left\{ -\frac{c}{\beta} \frac{k+1}{\tau_i + k} + \frac{k}{\tau_i + k - 1} \right\}$$

Now we can show that for any $\beta < c$ and for any i, the expression inside the brackets is negative, which will imply that the second-order derivative is negative, the function is concave, and the second-order condition is satisfied.

$$\frac{c}{\beta} \cdot \frac{k+1}{\tau_i + k} > \frac{k+1}{\tau_i + k} = \frac{k \qquad\qquad + 1}{(\tau_i + k - 1) + 1} > \frac{k}{\tau_i + k - 1}$$

■

8.4 Proof of Theorem 2

Recall that information types are indexed by their number of signals (τ_i) from most to least signals. Then we have that:

$$\frac{\tau_1 + k - 1}{\tau_1 + k} > ... > \frac{\tau_i + k - 1}{\tau_i + k} > ... > \frac{\tau_N + k - 1}{\tau_N + k}$$

Therefore:

$$\frac{p_1\frac{\tau_1}{\tau_1+k}}{p_1\frac{\tau_1}{\tau_1+k-1}} > ... > \frac{p_i\frac{\tau_i}{\tau_i+k}}{p_i\frac{\tau_i}{\tau_i+k-1}} > ... > \frac{p_N\frac{\tau_N}{\tau_N+k}}{p_N\frac{\tau_N}{\tau_N+k-1}}$$

Note that:

$$\frac{a}{b} > \frac{c}{d} \implies \frac{a}{b} > \frac{a+c}{b+d} > \frac{b}{d}$$

Therefore:

$$\frac{p_1\frac{\tau_1}{\tau_1+k}}{p_1\frac{\tau_1}{\tau_1+k-1}} > \frac{p_2\frac{\tau_2}{\tau_2+k}}{p_2\frac{\tau_2}{\tau_2+k-1}} \implies$$

$$\frac{p_1\frac{\tau_1}{\tau_1+k}}{p_1\frac{\tau_1}{\tau_1+k-1}} > \frac{p_1\frac{\tau_1}{\tau_1+k}+p_2\frac{\tau_2}{\tau_2+k}}{p_1\frac{\tau_1}{\tau_1+k-1}+p_2\frac{\tau_2}{\tau_2+k-1}} > \frac{p_2\frac{\tau_2}{\tau_2+k}}{p_2\frac{\tau_2}{\tau_2+k-1}}$$

We can then repeat this process with information type 3:

$$\frac{p_1\frac{\tau_1}{\tau_1+k}+p_2\frac{\tau_2}{\tau_2+k}}{p_1\frac{\tau_1}{\tau_1+k-1}+p_2\frac{\tau_2}{\tau_2+k-1}} > \frac{p_3\frac{\tau_3}{\tau_3+k}}{p_3\frac{\tau_3}{\tau_3+k-1}} \implies$$

$$\frac{p_1\frac{\tau_1}{\tau_1+k}+p_2\frac{\tau_2}{\tau_2+k}}{p_1\frac{\tau_1}{\tau_1+k-1}+p_2\frac{\tau_2}{\tau_2+k-1}} > \frac{p_1\frac{\tau_1}{\tau_1+k}+p_2\frac{\tau_2}{\tau_2+k}+p_3\frac{\tau_3}{\tau_3+k}}{p_1\frac{\tau_1}{\tau_1+k-1}+p_2\frac{\tau_2}{\tau_2+k-1}+p_3\frac{\tau_3}{\tau_3+k-1}} > \frac{p_3\frac{\tau_3}{\tau_3+k}}{p_3\frac{\tau_3}{\tau_3+k-1}}$$

Continuing until the final information type N we have that:

$$\frac{p_1\frac{\tau_1}{\tau_1+k}}{p_1\frac{\tau_1}{\tau_1+k-1}} > \frac{\sum_i p_i\frac{\tau_i}{\tau_i+k}}{\sum_i p_i\frac{\tau_i}{\tau_i+k-1}} > \frac{p_N\frac{\tau_N}{\tau_N+k}}{p_N\frac{\tau_N}{\tau_N+k-1}}$$

The middle expression is the ratio of expectations observed in the equilibrium suggestion threshold. With this realization, we can multiply all expressions by the hiring cost c to obtain:

$$x^1_{FB} = \frac{\tau_1+k-1}{\tau_1+k}c > \frac{\mathbb{E}\left[\frac{\tau_i}{\tau_i+k}\right]}{\mathbb{E}\left[\frac{\tau_i}{\tau_i+k-1}\right]}c > \frac{\tau_N+k-1}{\tau_N+k}c = x^N_{FB} \tag{30}$$

Thus we have that the first-best threshold for the information type with the most signals is greater than the equilibrium threshold while the first-best threshold for the information type with the least signals is less than the equilibrium threshold. Now consider the ratio of the probability of hire in the second best compared to the first-best for a given information type i:

$$R_i := \frac{Pr^{EQ}_i(hire)}{Pr^i_{FB}(hire)} = \left(\frac{x^i_{FB}}{x^{EQ}}\right)^k$$

The denominator is the same for all information types and k is a positive power, therefore the ratio is monotone decreasing in i (information type 1 has the highest ratio). From (30)

we know that $x^1_{FB}/x_{EQ} > 1$ and $x^N_{FB}/x_{EQ} < 1$. Thus the information type with the most signals are hired at higher rates in the equilibrium than the first-best, while the information type with the least is hired less in the equilibrium than the first-best. Further, because the ratio is monotone decreasing in the number of signals, all types with a number of signals above some cut-off will be hired at higher rates, while all those below some cut-off will be hired at lower rates.■

8.5 Proofs for Comparative Statics

8.5.1 Proof of Proposition 3

Consider all information types but i, and denote τ^*_{-i} the threshold τ for these information types.

$$\beta^* = \frac{(1-p_i)*\frac{\tau^*_{-i}}{\tau^*_{-i}+k}+p_i*\frac{\tau_i}{\tau_i+k}}{(1-p_i)*\frac{\tau^*_{-i}}{\tau^*_{-i}+k-1}+p_i*\frac{\tau_i}{\tau_i+k-1}}$$

$$\frac{d\beta^*}{d\tau_i} = \frac{p_i}{(1-p_i)*\frac{\tau^*_{-i}}{\tau^*_{-i}+k-1}+p_i*\frac{\tau_i}{\tau_i+k-1}} * \frac{\partial\frac{\tau_i}{\tau_i+k}}{\partial\tau_i} - \frac{\left((1-p_i)*\frac{\tau^*_{-i}}{\tau^*_{-i}+k}+p_i*\frac{\tau_i}{\tau_i+k}\right)*p_i}{\left((1-p_i)*\frac{\tau^*_{-i}}{\tau^*_{-i}+k-1}+p_i*\frac{\tau_i}{\tau_i+k-1}\right)^2} * \frac{\partial\frac{\tau_i}{\tau_i+k-1}}{\partial\tau_i}$$

$$\propto \frac{\frac{\partial\frac{\tau_i}{\tau_i+k}}{\partial\tau_i}}{\frac{\partial\frac{\tau_i}{\tau_i+k-1}}{\partial\tau_i}} - x^*_{EQ}$$

$$= \frac{k*(\tau_i+k-1)}{(k-1)*(\tau_i+k)} \cdot x^i_{FB} - x^*_{EQ}$$

$$> x^i_{FB} - x^*_{EQ} > 0$$

since $\tau_i > \tau^*$ and $k*(\tau_i+k-1) > (k-1)*(\tau_i+k)$. This implies that $\frac{d\beta^*}{d\tau_i} > 0$. ■

8.5.2 Proof of Propositions 4 and 5

$$\beta^* = \frac{\mathbb{E}\left[\frac{\bar{a}^k_i\tau_i}{\tau_i+k}\right]}{\mathbb{E}\left[\frac{\bar{a}^k_i\tau_i}{\tau_i+k-1}\right]}c = \frac{\sum_i p_i\bar{a}^k_i \cdot \left[\frac{\tau_i}{\tau_i+k}\right]}{\sum_i p_i\bar{a}^k_i \cdot \left[\frac{\tau_i}{\tau_i+k-1}\right]}c$$

Let us define the following measure function over i: $\tilde{p}_i = p_i\bar{a}^k_i$. Then β^* is given by the formula for the homogeneous $\bar{a}$ case

$$\beta^*/c = \frac{\mathbb{E}_{\tilde{p}}\left[\frac{\tau_i}{\tau_i+k}\right]}{\mathbb{E}_{\tilde{p}}\left[\frac{\tau_i}{\tau_i+k-1}\right]}$$

Increase in p_i or $\bar{a}_i$ is equivalent to increase in $\tilde{p}_i$. Assume $\tilde{p}_i$ increased to $\tilde{p}'_i$. The following inequalities are equivalent.

$$\frac{\tilde{p}_{-i}\frac{\tau^*_{-i}}{\tau^*_{-i}+k} + \tilde{p}'_i\frac{\tau_i}{\tau_i+k}}{\tilde{p}_{-i}\frac{\tau^*_{-i}}{\tau^*_{-i}+k-1} + \tilde{p}'_i\frac{\tau_i}{\tau_i+k-1}} > \frac{\tilde{p}_{-i}\frac{\tau^*_{-i}}{\tau^*_{-i}+k} + \tilde{p}_i\frac{\tau_i}{\tau_i+k}}{\tilde{p}_{-i}\frac{\tau^*_{-i}}{\tau^*_{-i}+k-1} + \tilde{p}_i\frac{\tau_i}{\tau_i+k-1}}$$

$$(\tilde{p}'_i - \tilde{p}_i) * \frac{\tau_i}{\tau_i+k} * \frac{\tau^*_{-i}}{\tau^*_{-i}+k-1} > (\tilde{p}'_i - \tilde{p}_i) * \frac{\tau^*_{-i}}{\tau^*_{-i}+k} * \frac{\tau_i}{\tau_i+k-1}$$

$$(\tilde{p}'_i - \tilde{p}_i) * \frac{\frac{\tau_i}{\tau_i+k}}{\frac{\tau_i}{\tau_i+k-1}} > (\tilde{p}'_i - \tilde{p}_i) * \frac{\frac{\tau^*_{-i}}{\tau^*_{-i}+k}}{\frac{\tau^*_{-i}}{\tau^*_{-i}+k-1}}$$

$$(\tilde{p}'_i - \tilde{p}_i) * \left(\frac{\frac{\tau_i}{\tau_i+k}}{\frac{\tau_i}{\tau_i+k-1}} - \frac{\frac{\tau^*_{-i}}{\tau^*_{-i}+k}}{\frac{\tau^*_{-i}}{\tau^*_{-i}+k-1}}\right) > 0$$

$$(\tilde{p}'_i - \tilde{p}_i) * \left(\frac{\tau_i+k-1}{\tau_i+k} - \frac{\tau^*_{-i}+k-1}{\tau^*_{-i}+k}\right) > 0$$

$$(\tilde{p}'_i - \tilde{p}_i) * (\tau_i - \tau^*_{-i}) > 0$$

Since x^τ_{FB} is increasing in τ. Therefore an increase in p_i or $\bar{a}_i$ increases x^*_{EQ} iff $\tau_i > \tau^*_{-i}$, which is equivalent to $\tau_i > \tau^*$. ■

8.6 Qualitative Interviews

We conducted three recorded interviews with people in the recruiting industry. The interviews occurred in 2019 and lasted around one hour each. The interviews asked a series of open-ended questions about recruiting in general and did not focus on the refund contract in particular. Audio recordings are available but require approval from the interviewees prior to release.

One interview was with an early career recruiter who works for an external recruiting firm. One interview was with a mid-career headhunter who at the time managed his own external recruiting firm. One interview was with an early career internal recruiter and human resource professional. The interview with the internal recruiter did not yield much information about external recruiters so it is not discussed much in this paper.

Both the mid-career external headhunter and the early-career external recruiter stated

that they utilized a contract structure where they got a fee that was contingent on the candidate staying for a certain period of time. Here is the excerpt from the mid-career headhunter:

Jacob (Author): "Do you ever consider the probability of termination or probability of leaving or retention when you are choosing who to suggest? Because you are paid based on placement but is it contingent on them staying for a certain period of time?"

Headhunter: "Yes most of the time it is anywhere from a 30 day to 90 day replacement. Sometimes it is full replacement no cost, sometimes it is a refund. We as recruiters may get paid 16,000 for a fee. But we can't..you know. You just don't want to send that money. Or you just won't get paid until those 90 days are up."

Here is the excerpt from the early-career recruiter:

Jacob (Author): "Do you consider probability of termination, probability of separation or retention when you are considering someone?"

Recruiter: "Yes."

Jacob (Author): "What is the window you get paid for generally?"

Recruiter: "Generally it is 90 days. We get paid upfront. One of two things happen. We either have the next placement for free or we return the money."

Jacob (Author): "And that's if they leave for what reasons?"

Recruiter: "Any reason."

Jacob (Author): "Even if the company fires them?"

Recruiter: "Yes."

Jacob (Author): "Do you ever contest why the company fired them?"

Recruiter: "No because the amount of times; I do not know if in my two years if we have ever had a candidate fired in the first 3 months."

Jacob (Author): "It is usually because they left or something else like that?"

Recruiter: "But even then it has happened only a few times. If a candidate leaves after the 90 day mark there are times when we may provide a discount on the next placement. But we are not tied to it by any means"

Jacob (Author): "Do you have any idea why it is 90 days?"

Recruiter: "It is just our good faith policy to our clients who we work with to say that hey we trust the candidates we are putting in front of you."

Jacob (Author): "It seems like everyone is 90 days, so. Like of all the people I have talked to, it seems like everyone is 90 days. Is there a reason for that number?"

Recruiter: "I have never had a conversation about it, my understanding would likely just be that turnover in tech is just so high we wodont want it to be past 3 months, but 3 months is enough time for someone to get caught up to speed and get to work and get

connected."

Jacob (Author): "One month is too short?"

Recruiter: "One month is too short. Six months is... there is too many reasons things could fall apart."

Additionally, when asked why companies use recruiter, one reason the early-career recruiter gave is: "Another reason is essentially we are free to use unless success happens. It doesn't cost any money unless they want to hire someone."

8.7 Proof of General Risk Attitudes

Proof. The firm's payoff for a suggested candidate with productivity a is equal to $\max\{a, 0\} - c$, which is a convex function. Let us take a candidate (μ, σ), which the firm decides to suggest. Then any candidate to the upper right from this one, with $\mu' \geq \mu$ and $\sigma' \geq \sigma$ is also suggested since the firm's payoff function is increasing and convex, and therefore a FOSD shift and an MPS increase its expectation. Thus, the firm suggests everyone above some threshold $\tilde{\mu}_{FB}(\sigma)$ that is decreasing in σ. This means that the firm is risk-loving since it prefers candidates with a mean-preserving spread of productivity. ■

Proof. The recruiter suggests anyone with (μ, σ) such that

$$\alpha - \beta \Pr(a < \beta | \mu, \sigma) \geq \bar{u}$$

$$\Pr(a < \beta | \mu, \sigma) \leq \frac{\alpha - \bar{u}}{\beta} \leq q$$

$\Pr(a < \beta | \mu, \sigma)$ is the conditional CDF of a given μ, σ, which we denote $F_{(\mu,\sigma)}(a)$. Then we have that: $\beta \leq F^{-1}_{(\mu,\sigma)}(q)$. Then for a fixed σ, consider all candidates with $\mu = \tilde{\mu}_{EQ}(\sigma)$ and $\sigma' < \sigma$. Since the candidates are q-lower-tail-risk ordered and $\beta \leq F^{-1}_{(\mu,\sigma)}(q)$:

$$\Pr(a < \beta | \mu, \sigma') < \Pr(a < \beta | \mu, \sigma) < q$$

And the (μ, σ') candidate should also be suggested. All candidates with σ and $\mu' > \tilde{\mu}_{EQ}(\sigma)$ are also suggested since the FOSD implies

$$\Pr(a < \beta | \mu', \sigma) < \Pr(a < \beta | \mu, \sigma) < q$$

This together implies that if a candidate (μ, σ) is suggested then all candidates to the upper left, $(\mu' \geq \mu, \sigma' \leq \sigma)$ should also be suggested, which implies that $\tilde{\mu}_{EQ}(\sigma)$ is increasing in σ. This result, on the other hand, can be thought of as risk aversion because the recruiter who suggests a candidate (μ, σ) would also suggest any other candidate (μ', σ') that second-order

stochastically dominates (μ, σ) candidate. ■