Why the Rooney Rule Fumbles: Limitations of Interview-stage Diversity Interventions in Labor Markets

Setareh Farajollahzadeh
McGill University, Montreal, Canada, setareh.farajollahzadeh@mcgill.ca

Soonbong Lee
Yale School of Management, New Haven, CT, soonbong.lee@yale.edu

Vahideh Manshadi
Yale School of Management, New Haven, CT, vahideh.manshadi@yale.edu

Faidra Monachou
Yale School of Management, New Haven, CT, faidra.monachou@yale.edu

Many industries, including the NFL with the Rooney Rule and law firms with the Mansfield Rule, have adopted interview-stage diversity interventions requiring a minimum representation of disadvantaged groups in the interview set. However, the effectiveness of such policies remains inconclusive. In light of this, we develop a framework of a two-stage hiring process, where rational firms, with limited interview and hiring capacities, aim to maximize the match value of their hires. The labor market consists of two equally sized social groups, m and w, with identical ex-post match value distributions. Match values are revealed only post-interview, while interview decisions rely on partially informative pre-interview scores. Pre-interview scores are more informative for group m, while interviews reveal more for group w; as a result, if firms could interview all candidates, both groups would be equally hired. However, due to limited interview capacity and information asymmetry, we show that requiring equal representation in the interview stage does not translate into equal representation in the hiring outcome, even though interviews are more informative for group w. In certain regimes, with or without intervention, a firm may interview more group w candidates, but still hire fewer. At an individual level, we show that strong candidates from both groups benefit from the intervention as the candidate-level competition weakens. For borderline candidates, group w candidates gain at the expense of group m. To understand the impact of non-universal interview-stage interventions on the market, we study a model with two vertically-differentiated firms, where only the top firm adopts the intervention. We characterize the unique equilibrium and demonstrate potentially negative effects: we show that in certain regimes, the lower firm hires less group w candidates due to increased firm-level competition for them, and further find examples where overall fewer group w candidates are hired across the market. At an individual level, while superstar candidates in both groups benefit, surprisingly the impact on borderline candidates may reverse: the lower firm may replace borderline group w candidates with borderline group m candidates in its interview set, effectively reducing the hiring probability of those borderline group w candidates. Overall, our findings highlight challenges in diversifying the labor market at early hiring stages due to information asymmetry, filtering, and competition. Beyond our context, our natural framework of a market with two-stage hiring may be of independent interest.

1. Introduction

In 2003, in response to concerns over the low representation of African Americans in head coaching and management positions, the NFL adopted the *Rooney rule*, named after Dan Rooney, the late owner of the Pittsburgh Steelers (NFL Football Operations 2025). Initially, the policy required teams with a head coaching vacancy to *interview* at least one racial minority candidate before making a hire; it later expanded to include more positions and women candidates. Taking a play from the NFL, many leading organizations across industries, from tech (Frier 2022) to higher education (Wong 2016) to law firms (Lamarre 2024), adopted similar interview policies to increase diversity in their workforce. For example, the *Mansfield rule*, which requires law firms to consider at least 30% of candidates for leadership roles from underrepresented groups, have been adopted by many leading law firms (Diversity Lab 2023). Overall, in 2023, 56% of the S&P 500 companies reported having an interview-stage diversity policy similar to the Rooney rule for higher-level positions (Spencer Stuart 2023).

Despite its widespread adoption, the Rooney Rule has had little to no success in increasing diversity in coaching teams, law firms, or C-suites (Cecchi-Dimeglio 2022, Rider et al. 2023, Smith and Lasker 2023). In 2022, Brian Flores, a former Black coach of Miami Dolphins, made headlines by filing a racial discrimination lawsuit against the NFL and the New York Giants, alleging that *"his 'Rooney rule' interview with the Giants was for a job that was never really open to him"* (Li and Dasrath 2022). Two separate statistical analyses of data from 2015 to 2022 revealed that one out of every three interviews conducted by NFL teams was with a Black candidate, but only 8 Black candidates were hired for a total of 56 open positions. The most statistically significant factor differentiating those hired from those not hired was race (Paine 2022, Smith and Lasker 2023). For law firms, Cecchi-Dimeglio (2022) found that the Mansfield rule certification did not have a noticeable impact on improving diversity. These examples motivate our research question:

> *Why do interview-stage diversity policies like the Rooney Rule struggle to significantly improve representation in hiring outcomes?*

1.1. Our Contributions

Various social factors, such as systematic biases and structural barriers, potentially contribute to the limited success of such diversity interventions. While these factors are profoundly important, we abstract away from them to focus on two *operational* aspects of the hiring process: (i) the two-stage nature of hiring with interviews and (ii) competition across firms. Our goal is to understand how these aspects influence the success of (or lack thereof) interview-stage diversity interventions like the Rooney rule. To that end, we develop and analyze a stylized model of a labor market with a two-stage hiring process.[1] Below, we provide an overview of our framework and main findings.

[1] We remark that a few other papers develop stylized models to study the Rooney rule; we provide a detailed comparison in Section 1.2.

A Labor Market Model with Two-Stage Hiring and Statistical Discrimination. In our model (Sections 2 and 4.1), each firm is fully rational, has limited hiring capacity, and aims to maximize the *match value* of its hires. However, match values are a priori unknown, so the firm has an *interview* stage to reveal the match values of interviewees. Each firm has limited interview capacity. Thus, it first selects a subset of candidates for an interview, and then hires the best based on their interview performance. Before the interviews, each firm can only observe noisy signals of candidates' match values, termed *pre-interview scores*, which capture their observable features such as prior experience and academic performance. Consequently, the firm relies on pre-interview scores to decide which candidates to interview. This two-stage hiring process mirrors common hiring practices across various sectors.

The pool of candidates consists of two equal-sized social groups, m and w. Assuming an idealized world without inherent biases and structural barriers, we posit that the two groups have identical match value distributions. Consequently, if the firms could interview everyone in the population, both groups would be equally represented in the hiring outcome. However, firms have limited interview capacity and must rely on (partially informative) pre-interview scores to allocate the limited interview slots between the two social groups.

The only distinction between the two groups arises from differences in the *informativeness* of their pre-interview scores and interviews. Specifically, following the seminal framework of Phelps (1972), we assume that pre-interview scores are normally distributed, with a common zero mean. Variance, however, differs across the two social groups, with group m's pre-interview scores having higher variance than group w's. This assumption is motivated by empirical evidence supporting the *"greater male variability hypothesis"* (Machin and Pekkarinen 2008, Baye and Monseur 2016), which suggests some asymmetry in the informativeness of features—in our model, the pre-interview scores—across social groups (see Assumption 1 and the related discussion). If the firm had a one-stage hiring process (i.e., without interviews) and hired candidates solely based on their pre-interview scores, the single-firm setting of our framework would essentially reduce to the classic statistical discrimination problem by Phelps (1972). In this case, the differential informativeness between the two groups, stemming from the difference in their pre-interview scores' variances, would directly lead to a lower hiring rate for group w, even though the firm is perfectly rational.

In our two-stage hiring process, however, firms do not make hiring decisions based on noisy scores; instead, they base their hiring decision on the match value of candidates revealed during interviews. Specifically, at the interview-stage, a candidate's match value is drawn independently across firms, conditional on their pre-interview score. The conditional match value is normally distributed with mean equal to the candidate's pre-interview score and known variance. This represents that a candidate with a higher pre-interview score is more likely to reveal a higher match value, but each firm further uncovers its own idiosyncratic preferences during interviews.[2] The variance again differs across groups,

[2] As we elaborate in Section 2, the interview process resembles the consumer choice framework, where the decision-maker (the firm in our setting) derives utility from two components: an intrinsic value (the candidate's pre-interview score) and an idiosyncratic component.

with group w now having higher variance than group m, meaning that the interview process is more informative for group w than group m to ensure the identical ex-post match value distribution across the two groups (see Assumption 1). As mentioned before, if firms could interview *all* candidates, both groups would be hired at an equal rate. However, due to the limited interview capacity and the information asymmetry across groups, disparities begin to arise, both at the group and individual levels. In contrast to the classical framework of Phelps (1972), the two-stage nature of hiring process and the distinct informational advantages of the two groups—namely, group m's pre-interview score (resp., group w's interview) being more informative than that of the other group—make the direction of discrimination highly non-trivial and introduce different channels of discrimination.

Analysis of a Single-firm Model: Limitation of the Rooney Rule. We begin our analysis by focusing on a single-firm setting (Section 3). We first characterize the firm's optimal group-aware solution without any policy intervention and show that it takes an intuitive greedy form in both stages (see Section 1.2 for further discussion of our benchmark). That is, the firm interviews (resp., hires) candidates with scores (resp., match values) above an interview (resp., hiring) threshold. Both thresholds are determined endogenously and interdependent. With this structural results, we show that, without any intervention, group w candidates will be interviewed (and then hired) at a lower rate than group m, for sufficiently low interview capacity (Theorem 1). Inspired by real-life policies such as the Rooney rule, we then consider an interview-stage intervention where the firm requires a minimum representation of group w among the interviewees. Hereafter, we use the term "the Rooney rule" to refer to this class of interview-stage interventions.

Our first main result is that imposing equal representation at the interview stage does not translate into equal representation in hiring outcomes: Group w candidates end up being hired at a lower rate than group m (Proposition 2). Even more strikingly, we show that in certain regimes, the firm may organically interview more group w candidates than group m candidates—rendering the intervention ineffective—yet still hire fewer from group w (Theorem 1 and Figure 2). These findings might seem counter-intuitive since group w's interviews are more informative than group m and at the same time match values have identical distributions. However, the greedy nature of the selection strategy for interviews acts as a *filtering mechanism* that favors more informative pre-interview scores (see Figure 1). Notably, this disparity arises due to limited interview capacity: with unlimited interviews, group w's more informative interviews would counterbalance their less informative pre-interview scores, but under the limited interview capacity, group w is systematically disadvantaged.

To analyze how the firm optimally allocates its limited interview capacity, we reveal a novel capacity-information trade-off that the firm faces, which can be viewed through an "exploration"-"exploitation" lens (Section 3.2.2). On the one hand, the firm wishes to interview more group w candidates since their interviews are more informative ("exploration"). On the other hand, due to limited interview capacity, the firm also prefers "safer" candidates from group m, whose pre-interview scores are more predictive

of their match values ("exploitation"). At lower interview capacity, the firm conservatively interviews most candidates from group m, but as its interview capacity increases, the firm becomes more open to the group w. From a technical perspective, the endogeneity of the firm's optimal hiring threshold adds technical complexity to the optimization problem. Our proof uses novel geometric arguments to address these technical challenges, characterizing three curves of interview thresholds: the iso-interview curve, the iso-hiring curve, and the optimal interview threshold curve as a function of interview capacity (see Section 3.2.3 and Figure 2-(b)).

At an individual level, the Rooney rule might benefit or harm different subgroups of candidates (Proposition 3). Overall, the candidate-level competition weakens across the interviewees, due to the endogenous change of hiring thresholds. Thus, "strong" candidates from both groups (i.e., those with sufficiently high scores) benefit from the intervention, since their hiring probability increases. For "borderline" candidates (i.e., those with scores near the interview thresholds), the Rooney rule redistributes interview opportunities from group m candidates to group w candidates. As a result, only group w candidates benefit at the expense of group m.

Our results are aligned with empirical findings that document the limited effectiveness of the Rooney rule in practice (as motivated earlier). Thus, in contrast to several theoretical works (Kleinberg and Raghavan 2018, Emelianov et al. 2020, Celis et al. 2021) that predict a positive impact of the Rooney rule, our results offer a different perspective, highlighting the limitations of such interventions. A key distinction lies in our modeling of the intermediate interview stage, as prior works typically assume a one-stage hiring process (see Section 1.2 for further discussion). Our results suggest that incorporating this feature into the operation of the market not only reflects the real-world better but also uncovers new channels—purely rational and statistical—that contribute to eventual disparity in hiring outcomes.

Analysis of Two-firm Model: Downstream Effects of the Rooney Rule. In the real world, firms rarely hire in isolation; they *compete* with each other for talents from the same pool of candidates. At the same time, interventions such as the Rooney rule are not universally adopted or enforced.[3] Motivated by these observations, we study the impact of these interventions in a two-firm vertically differentiated setting (Section 4). Both firms share the same pool of candidates (that consists of equisized groups m and w); however, only the top firm adopts the Rooney rule.

The hiring process remains almost identical in this two-firm setting: each firm undergoes its own interview-stage process;[4] however, the lower-ranked firm now faces competition from the top firm because all candidates prefer the top firm. Incorporating competition has a profound impact on the strategy of the lower-ranked firm: we show that in equilibrium the lower-ranked firm's optimal interview

[3] For example, a public report (Diversity Lab 2023) by a Mansfield certifier for law firms reveals that, there were at most 175 large-sized (resp., 65 mid-sized) law firms in the U.S. that were Mansfield certified in 2023, out of 264 large (resp., 168,044 mid-sized) firms in total (U.S. Census Bureau 2025).

[4] We remind that while candidates' pre-interview scores are shared across firms, their conditional match value revealed upon interview is independently drawn across firms.

strategy is not necessarily greedy (see Proposition 4 and Figure 3).[5] The lower-ranked firm may forego interviewing "superstar" candidates that are highly likely to be hired by the top firm. Instead, the lower-ranked firm may include "borderline" candidates not interviewed by the top firm.

With these structural results, we investigate the impact of the top firm's adoption of the Rooney rule. As discussed before, under the Rooney rule, the top firm replaces borderline group m candidates in its interview set with borderline group w candidates. Unburdened by the competition from the top firm, the lower-ranked firm may now include these borderline group m candidates in its interview set. As a domino effect, borderline group w candidates at the lower-ranked firm lose their interview spots. At a group level, this may lead to the hiring of fewer group w candidates by the lower-ranked firm (Proposition 5). This loss in hiring by the lower-ranked firm is not necessarily offset by the gain in hiring at the top firm: we find examples where overall fewer group w candidates are hired across the market. This primarily stems from the following: adoption of the Rooney rule weakens (resp., strengthens) the *firm-level competition* for group m (resp. w).

At an individual level, similar to the single-firm setting, "superstar" candidates from both groups—that is, candidates with exceptional scores who are exclusively interviewed by the top firm—gain due to *weakened candidate-level competition* in the top firm. However, somewhat surprisingly, the impact on borderline candidates may reverse: borderline group w candidates who gained an interview spot at the top firm but lost one in the lower-ranked firm may have a lower hiring probability because they face a *stronger candidate-level competition*. On the other hand, group m candidates who lost their spot in the top firm but now are interviewed by the lower firm may face a weaker candidate-level competition and thus have a higher hiring probability (Proposition 7). In contexts where being hired (by either firm) is substantially preferred over not being hired, our finding implies that adopting the Rooney rule can hurt some individuals in group w.

Overall, our results have profound policy implications: *the non-universal nature of policies such as the Rooney rule may indeed hurt the very group it intends to help at a group and individual level.* More broadly, our findings on the impact of the Rooney rule (in isolation or at the market level) highlight the challenges of improving representation in hiring outcomes using interview-stage interventions—often viewed as a form of "soft" affirmative action to create equal opportunities. Even in the absence of inherent biases or structural barriers, the interplay between information asymmetry and operational factors—such as limited interview capacities and competition—may limit the positive effect of such interventions or even lead to negative consequences.

We conclude by highlighting that a byproduct of our work is developing a natural framework for labor markets with interviews and characterizing its equilibrium. Furthermore, our equilibrium characterization as well as our comparative analysis (with vs. without intervention) involve many intricacies

[5] Our structural results sharply deviate from those of Vohra and Yoder (2023), which studies a similar model (motivated by questions unrelated to the Rooney Rule), as we discuss in detail in Section 1.2.

largely due to the interdependency between the interview and hiring decisions across groups and firms. Our framework and analysis can serve as a useful basis for future research in the areas of matching with interviews (Manjunath and Morrill 2023, Vohra and Yoder 2023, Ashlagi et al. 2025) and fairness in hiring processes (Hu and Chen 2018, Kleinberg and Raghavan 2018, Emelianov et al. 2020).

1.2. Related Literature

Our work broadly contributes to multiple strands of related literature, including statistical discrimination, (algorithmic) fairness in operations, and interviews in matching markets.

Statistical discrimination. The economics literature has proposed two prominent theories of discrimination: taste-based discrimination (Becker 1957) and statistical discrimination (Arrow 1971, Phelps 1972). Our framework builds on the seminal work of Phelps (1972) on statistical discrimination theory, where discrimination arises from differential evaluation uncertainty between two equally skilled groups. Building on Phelps (1972), we study the impact of differential evaluation uncertainty in a labor market with competition and a two-stage hiring process.

Several recent works extend Phelps' model in different dimensions (Kannan et al. 2019, Emelianov et al. 2020, Garg et al. 2020, Baek and Makhdoumi 2023). Closest to ours are the work of Garg et al. (2020) and Kannan et al. (2019). In the context of college admissions, Garg et al. (2020) extend Phelps' model to multiple features, strategic students facing test costs, and school-level competition. Their model is one-stage, as the final admission decision depends solely on a composite noisy signal of a student's true skill. In contrast, our two-stage model bases the eventual hiring decision on the match values revealed during the interview, rather than relying on a noisy signal (pre-interview score). This two-stage nature allows us to model the distinct informational advantages of the two groups at each stage, making the overall direction of disadvantage nontrivial.

Kannan et al. (2019) considers a two-stage Gaussian model in which a college admits students based on test scores, and an employer hires graduates based on grades—both being noisy signals of true skill. However, in their model, the school and employer are distinct decision-makers. Importantly, the college in the first stage does not solve an optimization problem; rather, the focus of Kannan et al. (2019) is on the existence of admission rules that ensure a certain notion of fairness rather than on optimal decision-making. In contrast, our model features a fully rational firm that optimizes hiring decisions across both stages, introducing technical complexities due to the interdependence between these decisions. Thus, our focus is on discrimination arising from such rational decision-making.

There is also an extensive line of empirical and experimental work documenting (statistical) discrimination in labor markets. The meta-analysis of field experiments by Quillian et al. (2017) shows no change in racial discrimination in hiring over time. Comprehensive surveys on the empirical studies of statistical discrimination can be found in Fang and Moro (2011), Guryan and Charles (2013), and Onuchic (2022). Our work compliments this literature by providing a theoretical framework of

statistical discrimination that incorporates both interviews and competition, and by studying why interview-stage diversity interventions may have limited impact with or without firm-level competition.

Fairness in operations. Motivated by societal, ethical, and legal considerations, a growing body of work studies fairness in operational decisions arising in a broad set of applications ranging from retail (e.g., pricing and assortment planning (Chen et al. 2022, Cohen et al. 2022, Chen et al. 2023, Manshadi et al. 2023b)) to public policy (e.g., college admissions (Bonet et al. 2024, Larroucau et al. 2024, Sirolly et al. 2024)) to political domains (Flanigan et al. 2021, Garg et al. 2022, Agiza et al. 2024). Additionally, while previously fairness has been mainly studied in static settings (Bertsimas et al. 2011), several recent papers focus on dynamic settings and conceptualize new notions of fairness (Freeman et al. 2017, Gupta and Kamble 2021, Allouah et al. 2023, Manshadi et al. 2023a). In the specific context of hiring, prior work has studied long-term fairness in labor markets (Hu and Chen 2018), fairness in sequential search (Aminian et al. 2023, Salem and Gupta 2024, Hu et al. 2025), fairness in online labor markets (Monachou and Ashlagi 2019, Gonzalez-Cabello et al. 2024), and algorithmic bias in resume screening (Cowgill 2020). For a survey on the topic of algorithmic fairness in hiring, the reader may refer to Fabris et al. (2024).

In particular, the Rooney rule has recently attracted attention in the fairness literature (Kleinberg and Raghavan 2018, Emelianov et al. 2020, Celis et al. 2021, Komiyama and Noda 2024, Kim et al. 2025), mainly focusing on its positive impact.[6] Our work contributes to this line of research by being the first to study the Rooney Rule beyond a single firm in a market with firm-level competition (Section 4). By doing so, our model not only uncovers new mechanisms for potentially negative downstream effects of the Rooney Rule but also highlights that insights from single-firm models do not necessarily generalize to more realistic settings with competition. In addition to studying the competition which sets our work apart from the prior works, even our single-firm setting has fundamental differences from closely related papers. We elaborate on these differences next.

Kleinberg and Raghavan (2018) study a discrete hiring model where a committee selects candidates to maximize the total potential. One group's potential is correctly observed, while the other's is subject to a constant multiplicative bias; their potential is drawn from the same Pareto distribution. They find that measures like the Rooney rule, which require selecting at least one candidate from the disadvantaged group, improve both representation and utility compared to the group-unaware approach. Our work differs in key ways: First, we adopt a statistical discrimination approach instead of assuming a fixed multiplicative bias. In other words, in our model there is no inherent bias against any group; the only difference between the groups is the informativeness of their pre-interview scores and interview

[6] A notable exception is Fershtman and Pavan (2021), who examine potential negative effects of soft affirmative action. A key distinction from our paper lies in the hiring model: they consider "dynamic interviewing," where a firm sequentially decides whether to expand their evaluation pool or assess a candidate already in the pool. In contrast, we adopt "batch interviewing," where firms pre-select a set of candidates, conduct all interviews at once, and then make hiring decisions. This structural difference leads to distinct mechanisms through which soft affirmative action may fail.

process. Second, in Kleinberg and Raghavan (2018), the firm bases its one-stage decision on biased information. However, in our two-stage model, the firm hires candidates based on true (match) values. As such, the underlying channels for discrimination are different.

Emelianov et al. (2020) study a statistical discrimination model with Gaussian distributions. While their base model is one-stage, they also consider a two-stage model in which the firm interviews candidates based on a noisy signal before hiring them based on true skills. The variance of the noisy signal varies across groups. Similarly, we assume that the variances of pre-interview scores differ between the two groups. However, a key difference is that, in their work, the firm is agnostic to such differences across groups and thus takes a *group-unaware* approach. In contrast, we adopt a *group-aware* framework where the firm, facing a continuum of candidates, knows each group's variance parameters and optimizes hiring decisions accordingly. We further show that leveraging this knowledge improves both diversity and utility compared to a group-unaware solution (see Appendix EC.3.12). Given that estimating population-level first-order statistics, such as variances, is plausible in various applications, the group-aware benchmark serves as a natural comparison point. Moreover, we analyze diversity outcomes under the group-aware benchmark in the two-stage setting and provide new insight into the limitations of the Rooney rule as well as its individual-level impacts (see Section 3.3).

Finally, the recent work of Kim et al. (2025) studies a two-stage hiring model where a minority group faces both multiplicative bias (as in Kleinberg and Raghavan (2018)) and statistical discrimination (as in Emelianov et al. (2020)). They show that first-stage interventions (such as the Rooney rule) are more effective than those at the second stage. A key distinction between our work and theirs is the incorporation of capacity constraints. Kim et al. (2025) assume no such constraints at either stage, instead modeling each stage with a fixed cost per candidate. In many real-world applications, however, fixed capacities are a natural operational constraint. Moreover, these constraints fundamentally impact both the theoretical analysis and findings. For instance, in Kim et al. (2025), the (second-stage) hiring threshold is *exogenously* set by the unit cost of hiring. As a result, first-stage interventions in their model do not affect this threshold or candidate-level competition. By contrast, in our model, the hiring threshold is *endogenously* determined by the interview set due to fixed capacities. Consequently, the Rooney Rule influences not only interview selection but also the hiring threshold. Notably, this adjustment reduces candidate-level competition, thereby benefiting strong candidates (Proposition 3).

Interviews in matching markets. A growing body of work explores matching with interviews, including Kadam (2015), Lee and Schwarz (2017), Echenique et al. (2022), Manjunath and Morrill (2023), Vohra and Yoder (2023), and Ashlagi et al. (2025). The closest work to ours is Vohra and Yoder (2023), which studies a similar two-stage hiring model motivated by market design questions unrelated to the Rooney rule or fairness aspects. Their analysis relies on specific distributional assumptions, particularly the "increasing k-yields" property for conditional match values. While a few distributions (e.g., exponential) satisfy this assumption, the normal distribution—a workhorse in statistical

discrimination—does not, nor does the Gumbel distribution—a workhorse in consumer choice modeling (see Remark 2). Consequently, the structural results in Vohra and Yoder (2023) do not apply to our setting.[7] In fact, this assumption fundamentally impacts the structure of interview sets in equilibrium: while Vohra and Yoder (2023) prescribes that even the lower-ranked firm greedily interviews superstars, our characterization (Proposition 4) shows that the lower-ranked firm may forgo interviewing superstars due to firm-level competition (see Remark 1 for more details).

2. Model

In this section, we present our baseline model featuring a single *firm*. Below, we outline the key components of the firm's two-stage hiring process. Then, we specify key assumptions regarding differences between the two groups and define the benchmark model. Finally, we formally introduce a model incorporating the interview-stage intervention. A table of key notation can be found in Appendix EC.1.

Candidates. A unit mass of candidates seek to be hired by the firm. Each candidate is characterized by two observable attributes: their *pre-interview score* $a \in \mathbb{R}$ and their *social group* $i \in \{m, w\}$. The pre-interview score a (henceforth, "score") signals the candidate's observable abilities, capturing factors such as standardized test scores or prior experience. While informative, the score is an imperfect signal of a candidate's true value. We assume that the scores of candidates from group i follow a Normal distribution, $a \mid i \sim \mathcal{N}(0, \sigma_i^2)$, where σ_i^2 denotes the group-specific score variance.

Beyond their score, each candidate has a *match value* v, which is the candidate's true value from the firm's perspective. However, candidates' match values are *ex-ante* unknown to the firm. Instead, candidates must undergo an interview, during which their match value v is observed. For a candidate from group i with score a, the interview reveals the candidate's match value drawn from the conditional distribution $v \mid \{a, i\} \sim \mathcal{N}(a, \tau_i^2)$, where τ_i^2 is the group-dependent variance of conditional match values. This reflects that higher scores are associated with higher match values, but the interview also reveals an idiosyncratic component of the candidate's match value. An alternative way to interpret the outlined interview process is to express the match value as $v = a + \epsilon$, where $\epsilon \mid i \sim \mathcal{N}(0, \tau_i^2)$ represents an idiosyncratic *interview signal* observed only post-interview and independent of the score a. This resembles the consumer choice modeling framework where a consumer's utility consists of two component, an intrinsic value and an idiosyncratic noise (Ben-Akiva 1985). Note that the ex-post (unconditional) match value distribution of group i is given by $v \mid i \sim \mathcal{N}(0, \sigma_i^2 + \tau_i^2)$.

The Firm. The firm seeks to hire a mass $\Delta \in (0, 1)$ of candidates, with the goal of maximizing the total match values of the hired candidates. Since the firm does not know the match values prior to interviews, its hiring process unfolds in the following two stages.

[7] For the same reason, we also highlight that the proof techniques in Vohra and Yoder (2023) do not apply to our setting. Instead, to derive structural our result, we develop a novel meta-characterization (see Propositions EC.1 and EC.2) that accommodates both Gaussian distributions and the class of distributions considered in Vohra and Yoder (2023). As a byproduct, the greedy structure of the lower-ranked firm's strategy in Vohra and Yoder (2023) follows directly from our results (see Remark EC.1 and Appendix EC.5.1).

(i) Interviewing: The firm selects which candidates to interview solely based on their observable attributes (their score a and social group $i \in \{m, w\}$). Formally, for each group $i \in \{m, w\}$, the firm selects an interview set $A_i \subseteq \mathbb{R}$, such that only candidates from group i with scores $a \in A_i$ are interviewed. The firm observes their match values drawn from the conditional distribution $v \mid \{a, i\} \sim \mathcal{N}(a, \tau_i^2)$, as described earlier. We assume that the firm has an interview capacity of $C \in [\Delta, 1]$.

(ii) Hiring: After the firm observes the match values v of interviewees, to maximize its total match value, the firm greedily hires the candidates with the highest non-negative match values until it reaches its hiring capacity Δ. This naturally leads to a threshold-based hiring policy where interviewees with match values above a hiring threshold $s \geq 0$ are hired.

Social Groups. We consider an idealized setting where the two groups $\{m, w\}$ are *ex-post symmetric* in the sense that the two groups are *equally-sized* and *equally-skilled*. Specifically, the two groups have the same population mass, and their match value distributions are identical. However, the groups differ in the variances of their scores and interview signals. Formally, we make the following assumption:

ASSUMPTION 1 **(Post-interview Identical but Pre-interview Non-identical Groups)**.

(a) The two groups have identical ex-post match value distributions: $\sigma_m^2 + \tau_m^2 = \sigma_w^2 + \tau_w^2$.

(b) However, $\sigma_m > \sigma_w$ *and* $\tau_w > \tau_m$.

Assumption 1-(a) implies that, from the firm's perspective, both groups are inherently identical in terms of their match values. Hence, in an ideal scenario with *unlimited interview capacity*, the hiring decision would be based on the identical ex-post match value distributions, resulting in equal representation of the two groups (given their equal population mass).

However, with limited interview capacity, the firm must select whom to interview solely based on their scores and group identity. This brings us to Assumption 1-(b), which introduces an asymmetry in their pre-interview score distribution. Specifically, we assume in part (b) that group m's scores are *more informative* of their match value than those of group w ($\sigma_m > \sigma_w$). This is in line with empirical evidence showing differences in the informativeness of pre-interview measures across social groups. For example, Machin and Pekkarinen (2008) and Baye and Monseur (2016) support the *"greater male variability hypothesis,"* showing a higher representation of males in the tails of test score distributions. Furthermore, Rothstein (2004) shows that SAT scores are more informative for high-income than low-income students. Additionally, this assumption is aligned with the literature on statistical discrimination (in one-stage models) once we view our score as the counterpart of the "skill estimate" that the firm forms, based on which it makes a decision. For instance, building on Phelps (1972), Garg et al. (2020) show that the differential informativeness in the features across social groups leads to lower variance in the skill estimates of the minority group (see their Lemma 1 and Figure 1).

Note that, due to the post-interview symmetry implied by part (a), group m's higher informativeness in their scores must be coupled with a *less* informative interview ($\tau_m < \tau_w$). That is, each group has

a distinct informational advantage: group m's pre-interview score is more informative, whereas group w reveals more information during the interview.[8] Given this non-trivial information asymmetry and limited interview capacity, the rational firm carefully optimizes its interview and hiring decision to maximize the total match value of the hired candidates.

Firm's Optimization Problem (Benchmark). We now formally state the firm's optimization problem. For each group i, let $H_i(a)$ and $h_i(a)$ denote the cumulative distribution function (C.D.F.) and probability density function (p.d.f.) of their respective score distributions. Similarly, let $G_i(v \mid a)$ and $g_i(v \mid a)$ denote the C.D.F. and p.d.f. of the conditional match value distribution, with $\bar{G}_i(v \mid a) = 1 - G_i(v \mid a)$ denoting the complementary C.D.F. As highlighted earlier, the firm selects an interview set based on the candidate's score a and group identity i, which we denote by $A = (A_m, A_w)$. The firm further selects a non-negative hiring threshold s to maximize the total match value of the hired candidates. Formally, the firm solves the following optimization problem:

$$\text{(Benchmark)} \quad \max_{\substack{A_m, A_w \subseteq \mathbb{R} \\ s \geq 0}} \sum_{i \in \{m,w\}} 0.5 \int_{A_i} \int_{s}^{\infty} v g_i(v \mid a) h_i(a) \, \mathrm{d}v \, \mathrm{d}a \tag{1}$$

$$\text{s.t.} \sum_{i \in \{m,w\}} 0.5 \int_{A_i} h_i(a) \, \mathrm{d}a = C \tag{2}$$

$$\sum_{i \in \{m,w\}} 0.5 \int_{A_i} \bar{G}_i(s \mid a) h_i(a) \, \mathrm{d}a \leq \Delta \tag{3}$$

We refer to this optimization problem as (Benchmark). The objective function (1) represents the total match value of the hired candidates.[9] Note that, even in a group-aware framework, the firm maximizes its overall match value by setting the same hiring threshold for the two groups. Constraint (2) ensures that the interview mass equals interview capacity C.[10] Constraint (3) ensures that the hiring mass does not exceed hiring capacity Δ. We emphasize that the hiring capacity constraint (3) is expressed as an inequality. This choice is intentional: it may not always be optimal to fully fill the hiring capacity Δ since the match values can be negative (see Proposition 1).

We use $A^* = (A_m^*, A_w^*)$ and s^* to denote the optimal solution of (Benchmark). We further define ρ^* and π^* as the optimal interview and hiring fractions of group w, respectively:

$$\rho^* := \frac{0.5 \int_{A_w^*}^{\infty} h_w(a) \, \mathrm{d}a}{C}, \quad \pi^* := \frac{\int_{A_w^*} \bar{G}_w(s^* \mid a) h_w(a) \, \mathrm{d}a}{\sum_{i \in \{m,w\}} \int_{A_i^*} \bar{G}(s^* \mid a) h_i(a) \, \mathrm{d}a}. \tag{4}$$

[8] In reality, one group's interview may not necessarily be more informative than the other's. However, we emphasize that our assumption, $\tau_m < \tau_w$, is necessary to capture an *idealized* setting where there is no inherent difference in the true match values between the two groups. If instead $\tau_m = \tau_w$, hiring disparities persist *even with unlimited interview capacity*. Specifically, when $\Delta < 0.5$, the group w is consistently hired less than group m even if $C = 1$.

[9] Because the firm faces a continuum (i.e., infinite) candidates, we implicitly use the law of large numbers to formulate its optimization problem as deterministic—see Proposition 2.1 of Emelianov et al. (2020) for technical details.

[10] It is straightforward to show that, even if the constraint (2) is expressed as a weak inequality, it remains optimal for the firm to fully utilize the interview capacity C.

That is, if $\rho^* < 0.5$ (resp., $\pi^* < 0.5$), group w is under-represented in the interview-stage (resp., final hiring outcome) under the optimal solution of (Benchmark).

Interview-stage Intervention (ρ-Rooney rule). Motivated by potential disparities arising in the benchmark and policies like the Rooney rule, we consider an *interview-stage* intervention that the firm may adopt to increase the representation of group w in the final hiring outcome. Concretely, we formalize this intervention as the *ρ-Rooney rule*, which requires that the fraction of group w candidates in the interview set must be *at least* $\rho \in (0,1)$:

$$0.5 \int_{A_w} h_w(a)\,\mathrm{d}a \geq \rho C \qquad (\rho\text{-Rooney rule})$$

Hereafter, we interchangeably use the ρ-Rooney rule and the interview-stage intervention. When the firm adopts ρ-Rooney rule, with the same two-stage hiring process outlined earlier, the firm solves the following optimization problem:

$$\text{(Intervention)} \quad \max_{\substack{A_m, A_w \subseteq \mathbb{R} \\ s \geq 0}} \sum_{i \in \{m,w\}} 0.5 \int_{A_i} \int_{s}^{\infty} v g_i(v \mid a) h_i(a)\,\mathrm{d}v\,\mathrm{d}a \quad \text{s.t.} \quad (2), (3), (\rho\text{-Rooney rule})$$

We often refer to this optimization problem as (Intervention). We use $A^\rho = (A^\rho_m, A^\rho_w)$ and s^ρ to denote the optimal solution under the ρ-Rooney Rule.

3. Analysis of the Single-Firm Model

In this section, we examine the impact of the ρ-Rooney Rule in the single-firm model. We first characterize the firm's optimal decision in Section 3.1, then explore the group-level and individual implications of the ρ-Rooney rule in Section 3.2 and Section 3.3, respectively.

3.1. Structural Results

We begin by characterizing the firm's optimal interview set and hiring threshold under (Benchmark). In the following proposition, we show that the firm's optimal decisions for interview and hiring have a greedy structure. (For $x \in \mathbb{R}$, we use x_+ to denote a positive part of x, i.e., $x_+ = \max\{x, 0\}$.)

PROPOSITION 1 **(Characterization of Benchmark Optimal Solution).** *The firm's optimal interview set $A^* = (A^*_m, A^*_w)$ and hiring threshold s^* uniquely exist and are jointly determined as follows:*

(a) ***Greedy Optimal Interview Set:*** *Let $F_i(a, s^*) := \mathbb{E}[(v - s^*)_+ \mid a, i]$. Then, the firm's optimal interview set A^*_i for each group i is a super-level set of $F_i(a, s^*)$, i.e.,*

$$A^*_i = \{a \in \mathbb{R} : F_i(a, s^*) \geq \theta\}, \tag{5}$$

where level θ is chosen to satisfy the interview capacity constraint (2). Furthermore, there exists interview threshold $a^*_i \in \mathbb{R}$ such that $A^*_i = [a^*_i, \infty)$.

(b) ***Greedy Optimal Hiring Threshold:*** *The optimal hiring threshold s^* satisfies*

$$s^* = \min\Big\{s \geq 0 : \sum_{i \in \{m,w\}} 0.5 \int_{A^*_i} \bar{G}_i(s \mid a) h_i(a)\,\mathrm{d}a \leq \Delta\Big\}. \tag{6}$$

Proposition 1 establishes that the optimal interview set and hiring threshold uniquely exist and exhibit interdependent greedy structures. Specifically, within each group, the firm selects candidates to interview greedily based on their pre-interview scores (part (a)). Consequently, the firm only needs to determine the *interview thresholds* $\boldsymbol{a}^* = (a_m^*, a_w^*)$ for each group — see Figure 1-(a) for an illustration of this greedy interview set. Post-interview, the firm's hiring strategy is similarly greedy: the firm hires the interviewees with the highest non-negative match values v until it fills its hiring capacity (part (b)). Notably, due to the capacity constraints, the optimal interview sets and hiring threshold jointly depend on each other (Equations (5) and (6)), adding technical complexity to finding the optimal solution.

The greedy structure of the interview set is driven by the function $F_i(a, s^*) = \mathbb{E}[(v - s^*)_+ \mid a, i]$, referred to as the *excess value* function, which represents the expected additional match value of a candidate from group i with score a, *above the optimal hiring threshold* s^*. Intuitively, it captures how valuable a candidate is to the firm relative to the hiring threshold. As such, the firm orders candidates based on their excess value and interviews all candidates above common level θ, regardless of their group (see Equation (5)). The high-score candidates naturally have greater potential to reveal higher match values, and therefore $F_i(a, s^*)$ increases with a (we formally prove this in Appendix EC.3.2). Thus, the optimal interview set in Equation (5) is greedy in the score a. Furthermore, given the optimal hiring threshold s^*, Equation (5) uniquely determines the interview set A_i^* (and the corresponding interview threshold a_i^*) by selecting the level θ that satisfies the interview capacity constraint (2).

We prove Proposition 1 in Appendix EC.2 and EC.3.2. First, we establish Equations (5) and (6) as necessary optimality conditions through proof by contradiction. Specifically, we show that if the pair of optimal interview set (A_m^*, A_w^*) and hiring threshold s^* fails to satisfy these conditions, we can construct an alternative feasible pair that strictly improves the firm's objective (1). To do so, we use an intricate exchange argument that replaces certain candidates in the interview set with others who yield a higher excess value, while adjusting s^* accordingly to respect the hiring capacity. Next, we establish uniqueness by showing that a solution satisfying the optimality conditions (5) and (6) uniquely exists. (Note that this further implies that Equations (5) and (6) are also sufficient conditions for optimality.) If the hiring threshold s^* were known, the interview set (A_m^*, A_w^*) would be uniquely determined by Equation (5). However, s^* itself depends on (A_m^*, A_w^*) through Equation (6). Based on these observations, we reformulate the system of Equations (5) and (6) as a fixed point equation for s^*. In Proposition EC.2, we show that such a fixed point uniquely exists. Finally, in Appendix EC.3.2, we establish that the excess value $F_i(a, s^*)$ increases in a using the notion of likelihood ratio order (Shaked and Shanthikumar 2007), which directly leads to the greedy structure of A_i^*.

While the firm's optimal interview set is greedy in scores within each group, the interview thresholds a_m^* and a_w^* may differ. This difference arises because conditional on the same score, group w candidates' match values have higher variance than group m (Assumption 1). Thus, the rational firm must account for group w's potential to reveal more extreme match values during interviews. Formally, we show that $a_w^* < a_m^*$ in the following corollary, which we prove in Appendix EC.3.3.

COROLLARY 1 **(Lower Optimal Interview Threshold for Group w than Group m).** *The optimal interview threshold for group w is strictly lower than that for group m, i.e., $a_w^* < a_m^*$.*

We prove Corollary 1 in Appendix EC.3.3. The proof leverages the optimality condition (5), which establishes an intuitive relationship between the interview thresholds: the excess value must be equal at the interview thresholds for both groups, i.e.,

$$F_m(a_m^*, s^*) = F_w(a_w^*, s^*). \tag{7}$$

We refer to this equation as the *balance condition*. Using the fact that $\tau_w > \tau_m$ (Assumption 1), we show that this condition can only hold when $a_m^* > a_w^*$ through the notion of increasing convex order (Shaked and Shanthikumar 2007) (see Appendix EC.3.1 for a formal definition).

We conclude this section by characterizing the optimal solution under the ρ-Rooney rule. If the optimal solution under (Benchmark) satisfies the ρ-Rooney rule (i.e., $\rho \le \rho^*$; see Equation (4)), the firm's solution remains unchanged. Otherwise, the following corollary, proven in Appendix EC.3.4, shows that the firm adjusts the interview thresholds to ensure that group w's interview fraction equals ρ. Specifically, the firm lowers the interview threshold for group w while raises it for group m.

COROLLARY 2 **(Characterizing the Optimal Solution under the ρ-Rooney rule).** *For any $\rho \in [0,1]$, the firm's optimal interview set $A^\rho = (A_m^\rho, A_w^\rho)$ and hiring threshold s^ρ under the ρ-Rooney rule uniquely exist and are jointly determined as follows:*

(a) The optimal interview set for group i is given by $A_i^\rho = [a_i^\rho, \infty)$, where a_w^ρ is given by

$$a_w^\rho = \begin{cases} a_w^* & \text{if } \rho \le \rho^* \\ H_w^{-1}(1 - 2\rho C) & \text{otherwise,} \end{cases} \tag{8}$$

and a_m^ρ is uniquely determined by the interview capacity constraint (2). Furthermore, we have $a_m^\rho > a_m^*$ and $a_w^\rho < a_w^*$ whenever $\rho > \rho^*$.

(b) The optimal hiring threshold s^ρ satisfies:

$$s^\rho = \min\left\{ s \ge 0 : \sum_{i \in \{m,w\}} 0.5 \int_{A_i^\rho} \bar{G}_i(s \mid a) h_i(a)\, \mathrm{d}a \le \Delta \right\}. \tag{9}$$

3.2. Group-level Implications of Interview-Stage Intervention

Equipped with the structural characterization in Proposition 1 and its corollaries, we now analyze impacts of (ρ-Rooney rule) on group-level hiring outcomes. Our main focus is the 0.5-Rooney Rule, which represents the most natural and ideal form of the ρ-Rooney Rule given the equal population mass of the two groups. However, our qualitative insights extend to general values of ρ.

3.2.1. Equal Interview Representation

To build intuition, suppose for a moment that the firm uses a greedy interview set (similar to Proposition 1 and Corollary 2) but requires the two groups to be interviewed at the same rate. Specifically, for any limited interview capacity $C < 1$, the firm follows the interview threshold $\boldsymbol{a}' = (a'_m, a'_w)$ such that the interview mass of both groups equals $0.5C$. Would such a policy lead to equal representation in the hiring outcome? If not, which group would face a disadvantage? A priori, the answer is not clear since the two groups have distinct informational advantages: compared to group m, group w's scores are less informative, whereas their interview signals are more informative about their match value.

Note that, under the equal interview rates, comparing the hiring masses of the two groups reduces to comparing their conditional distributions of match value v given that a candidate is interviewed (i.e., $v \mid \{a \geq a'_i, i\}$). To do so, it is instructive to study the group i's posterior distribution of score a given match value v. Since $a \mid i \sim \mathcal{N}(0, \sigma_i^2)$ and $v \mid \{a, i\} \sim \mathcal{N}(a, \tau_i^2)$, the posterior score distribution is $a \mid \{v, i\} \sim \mathcal{N}\left(\frac{\sigma_i^2}{\sigma_i^2+\tau_i^2} v, \left(\frac{1}{\sigma_i^2} + \frac{1}{\tau_i^2}\right)^{-1}\right)$. Therefore, the posterior mean of a group i candidate's score given match value v, normalized by its posterior standard deviation, is given by:

$$\frac{\mathbb{E}[a \mid v, i]}{\sqrt{\mathrm{Var}[a \mid v, i]}} = \frac{\sigma_i}{\tau_i} \frac{v}{\sqrt{\sigma_i^2 + \tau_i^2}}. \tag{10}$$

Equation (10) suggests the greater power of pre-interview scores compared to interview signals in this two-stage hiring process. Specifically, given Assumption 1, the normalized posterior mean of scores in Equation (10) is higher for group m than for group w when v is positive. Thus, group m candidates with high v tend to have higher pre-interview scores on average, relative to their variance, compared to their group w counterparts. This creates an informational advantage for group m.

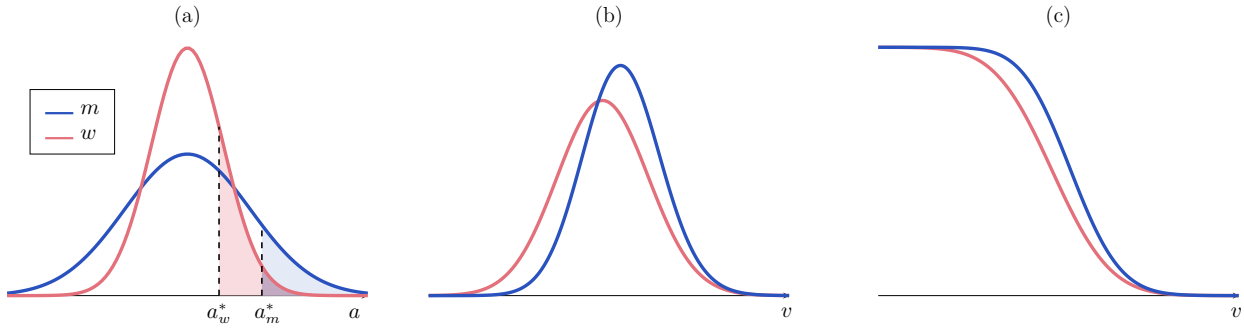


Figure 1 **Panel (a): the firm's optimal interview strategy (greedy in the score for each group; see Proposition 1). Panel (b): Density of $v \mid \{a \geq a'_i, i\}$ where the interview threshold $\boldsymbol{a}' = (a'_m, a'_w)$ equalizes the interview mass for both groups. Panel (c): Complementary C.D.F of the same distribution from panel (b) (see Proposition 2).**

If the firm did not face an interview capacity constraint, group w's more informative interview signals would ultimately counterbalance group m's informational advantage in pre-interview scores. However, under limited interview capacity (thus the interview threshold $\boldsymbol{a}'$ being finite), the firm's

greedy interview structure, which prioritizes candidates with higher scores, acts as a *filtering mechanism* that leverages group m's advantage. Specifically, it makes group m candidates with positive v more likely to enter the interview set compared to their group w counterparts. Consequently, due to the equal interview rates, group m's conditional distribution of v given being interviewed, $v \mid \{a \geq a'_m, m\}$, shifts to the right relative to group w's, $v \mid \{a \geq a'_w, w\}$ as we illustrate in Figure 1-(b).

This rightward shift results in the stochastic dominance of the match value distribution for group m interviewees over that of group w interviewees. Specifically, due to the higher normalized posterior mean of scores in Equation (10) for group m and the fact that both groups are interviewed at equal rates, the density function of $v \mid \{a \geq a'_m, m\}$ intersects that of $v \mid \{a \geq a'_w, w\}$ *only once*, with group m having higher density for larger v—see Figure 1-(b). This single-crossing property leads to the stochastic dominance of $v \mid \{a \geq a'_m, m\}$ over $v \mid \{a \geq a'_w, w\}$, as confirmed in Figure 1-(c) via the complementary C.D.F.s of the two distributions. Consequently, for any finite hiring threshold s, the right tail of $v \mid \{a \geq a'_m, m\}$ has greater mass than that of $v \mid \{a \geq a'_w, w\}$. This implies that, under equal interview rates, group m is hired at a higher rate than group w. We summarize and formalize this intuition in the following proposition, which we prove in Appendix EC.3.5.

Proposition 2 **(Group w is Under-represented in Hiring under Equal Interview Rates).** *For any limited interview capacity $C < 1$ and interview threshold $\boldsymbol{a}' = (a'_m, a'_w)$ that equalizes the interview mass of the two groups, the resulting hiring mass of group w is strictly less than that of group m for any finite hiring threshold s. Formally, if the finite interview threshold $\boldsymbol{a}' = (a'_m, a'_w)$ satisfies*

$$\int_{a'_m}^{\infty} h_m(a)\,\mathrm{d}a = \int_{a'_w}^{\infty} h_w(a)\,\mathrm{d}a, \tag{11}$$

then, for any finite s, we have

$$\int_{a'_m}^{\infty} \bar{G}_m(s \mid a) h_m(a)\,\mathrm{d}a > \int_{a'_w}^{\infty} \bar{G}_w(s \mid a) h_w(a)\,\mathrm{d}a. \tag{12}$$

Proposition 2 highlights the limited impact of the 0.5-Rooney rule: even if the firm ensures demographic parity at the interview stage, group w remains under-represented in the final hiring outcome. This result suggests that the hiring disparity does not stem from the unequal representation at the interview stage, but rather from structural differences in the informativeness of candidate's signals and limited interview capacity. Although in our idealized setting interview signals favor group w, the firm's greedy interviewing strategy amplifies group m's pre-interview advantage, perpetuating the disparity from the pre-interview stage to the final hiring outcome.

3.2.2. Optimal Interview Representation

In its optimal decision-making, the firm does not necessarily interview the two groups equally. Specifically, the 0.5-Rooney rule can meaningfully alter the firm's decisions only if (Benchmark) violates the rule (see Corollary 2). Otherwise, if the firm already finds it beneficial to comply with the rule, the

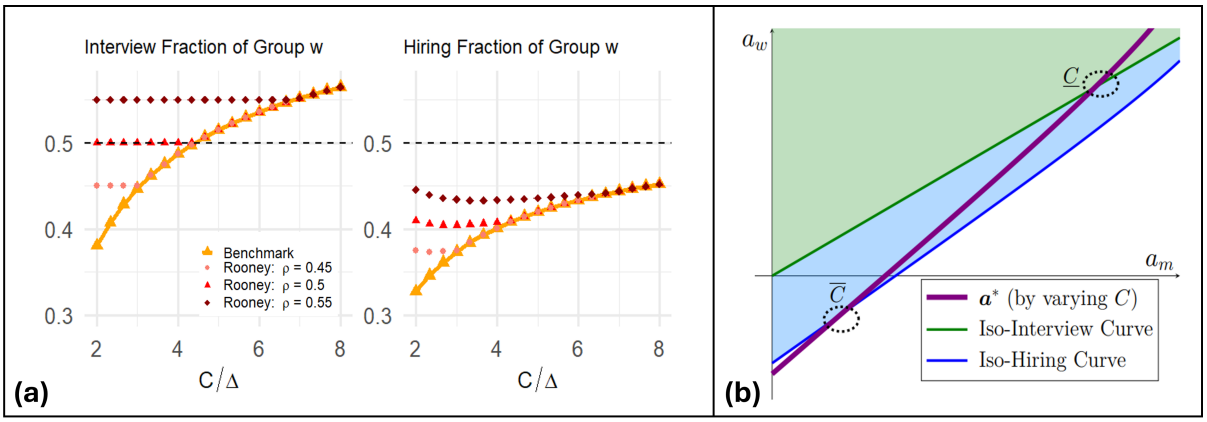


Figure 2 **Panel (a): the interview and hiring fraction of group w (Theorem 1). We vary interview capacity C while fixing other parameters Δ=0.01 and $(\sigma_m^2, \sigma_w^2) = (\tau_w^2, \tau_m^2) = (6, 4)$, with the horizontal axis representing the normalized interview capacity (C/Δ). Panel (b): proof sketch of Theorem 1.**

intervention has no impact on hiring outcomes. Thus, the key to assessing the 0.5-Rooney rule's effect is understanding the firm's *optimal interview representation* of group w under (Benchmark).

We find that the representation of group w under (Benchmark) depends on the interview capacity C. Specifically, the firm faces a trade-off in allocating its limited interview slots between the two groups. On the one hand, compared to group m, group w candidates have less informative scores. On the other hand, group w candidates reveal more information during interviews than group m. This can be seen through an "exploration"-"exploitation" lens. Interviewing group w candidates involves "exploration"—taking chances on candidates whose high match values might be revealed only through the interview. In contrast, interviewing group m candidates represents "exploitation"—selecting "safe" candidates whose scores already provide strong signals about their match values. When C is small (relative to Δ), the firm must allocate its limited interview slots carefully, making the cost of exploration higher. Consequently, the firm prioritizes the safer, high-score group m candidates, leading to group w's under-representation in both the interview set and the hiring outcome. To illustrate this intuition, Figure 2-(a) shows the interview and hiring fraction of group w candidates under (Benchmark) (orange line) by varying C for a fixed Δ. In the figure, we observe that for small values of C (for example, $C/\Delta \leq 4$), the firm interviews (and hires) fewer group w candidates compared to group m.

As C increases, the firm gains more capacity to explore, allowing it to interview more group w candidates. This can even result in group w's over-representation in the interview set. However, as highlighted earlier, the firm's greedy interview structure favors group m in the hiring outcome. Consequently, and perhaps counterintuitively, for moderate values of C—where the firm focuses on "high-enough" score candidates— group w may remain under-represented in the hiring outcomes despite being over-represented in the interview set. Figure 2-(a) again confirms this intuition: for larger values of C (for example, $C/\Delta > 4$), the firm interviews more group w but still hire fewer than group m.

Impact of the 0.5-Rooney rule. In summary, when interview capacity C is small enough, the firm under (Benchmark) under-represents group w in the interview set. As such, by Corollary 2, the 0.5-Rooney rule forces the firm to adjust its interview sets to impose equal interview rates, thereby increasing group w's interview mass. Naturally, this adjustment increases the group's hiring mass. However, as established in Proposition 2, it *does not* lead to equal representation in hiring (see red triangles in Figure 2). When C is larger, the firm *organically* over-represents group w in interviews but may still under-represent them in the final hiring outcome due to the filtering mechanism of greedy interviews (as explained in Section 3.2.1). In this case, the 0.5-Rooney rule has no meaningful impact on hiring outcomes, as (Benchmark) already satisfies the 0.5 Rooney rule — as shown in Figure 2.

We formalize all the above observations through the following theorem.

THEOREM 1 **(Limitations of the 0.5-Rooney rule).** *For small enough Δ, there exist interview capacity thresholds $\underline{C} < \overline{C}$ such that, under (Benchmark),*

(a) ***The 0.5-Rooney rule has non-zero but limited effect****: If $C < \underline{C}$, group w is under-represented in both of the interview set and hiring outcome (i.e., $\rho^* < 0.5$ and $\pi^* < 0.5$ where ρ^* and π^* are defined in Equation (4)). In this case, the 0.5-Rooney rule strictly increases group w's hiring mass, but the group still remains under-represented in the hiring outcome.*

(b) ***The 0.5-Rooney rule has no effect****: If $C \in [\underline{C}, \overline{C})$, group w is over-represented in the interview set (i.e., $\rho^* \geq 0.5$) but under-represented in the hiring outcome (i.e., $\pi^* < 0.5$). In this case, the 0.5-Rooney rule has no effect.*

We outline the proof of Theorem 1 in the next subsection. Although Theorem 1 focuses on the 0.5-Rooney Rule, our simulation suggests that its insights naturally extend to general ρ. For $\rho \leq 0.5$, group w's hiring fraction falls short of ρ when the firm adopts the ρ-Rooney rule. Figure 2-(a) illustrates this for $\rho = 0.45$. Conversely, as implied by Proposition 2—particularly the strict inequality (12)—a value of ρ above 0.5 may still fail to achieve equal hiring. This is because the only way to increase group w's hiring mass given a fixed interview capacity is to decrease (resp., increase) the interview threshold of group w (resp., group m) — see Lemma EC.3 in Appendix EC.3.6. As an illustrative example, Figure 2-(a) shows that even with $\rho = 0.55$, group w remains underrepresented in the hiring outcomes.

3.2.3. Proof Sketch of Theorem 1

In this section, we briefly sketch the proof of Theorem 1. A complete proof with a more detailed outline can be found in Appendix EC.3.6. Consider a scenario where the firm follows a greedy strategy for interviewing with an interview threshold vector $\boldsymbol{a} = (a_m, a_w)$. In Figure 2-(b), fixing a sufficiently small Δ, we display three different sets of the interview thresholds for group m (x-axis) and group w (y-axis). The green curve represents the *iso-interview* curve, the set of interview thresholds where the two groups are interviewed at equal rates (i.e., that satisfies Equation (11)). On the other hand, for a fixed a_m, we show that there exists a unique value of a_w that equates the hiring mass of the two

groups, when the hiring threshold is optimally chosen given these exogenous interview thresholds. The set of such interview thresholds, which we refer to as the *iso-hiring* curve, is displayed as the blue curve in Figure 2-(b). See Definition EC.6 and Proposition EC.3 for a precise definition of the iso-interview and iso-hiring curves. Intuitively, fixing group m's interview threshold, group w's interview and hiring mass decrease as its interview threshold increases (see Lemma EC.3). Hence, the epigraph of the iso-interview curve (resp., iso-hiring curve) represents the set of interview thresholds where group w is under-represented in the interview set (resp., hiring outcomes).[11] Finally, the purple curve represents the optimal interview thresholds $\boldsymbol{a}^*$ under (Benchmark) as the interview capacity $C \geq \Delta$ varies. The curve moves downward and to the left as C increases, meaning that the firm interviews more candidates from both groups by lowering their thresholds (we formally establish this in Appendix EC.3.9).

The crux of our proof is to examine how the curve of optimal interview thresholds intersects the iso-interview and iso-hiring curves as the interview capacity C increases. Specifically, we build on Figure 2-(b) to prove Theorem 1 in three steps. In Step 1, we show that the iso-interview curve lies above the iso-hiring curve (Proposition EC.3), as illustrated by the relative positions of the blue and green curves in Figure 2-(b). In Step 2, we show that the curve of optimal interview thresholds intersects the iso-interview curve only once (Proposition EC.4). That is, (Benchmark) over-represents group w in the interview set if and only if the interview capacity C exceeds some threshold $\underline{C}$ — see the intersection of purple and green curves in Figure 2-(b). However, since $\boldsymbol{a}^*$ varies continuously with C (formally established in Appendix EC.3.9), the optimal interview threshold can intersect the iso-hiring curve only when the interview capacity reaches some *higher* threshold $\overline{C} > \underline{C}$. In the final step, we analyze the optimal interview and hiring representation of group w across two ranges of C: (i) $C < \underline{C}$ (green-colored area in Figure 2-(b)), and (ii) $C \in [\underline{C}, \overline{C})$ (blue-colored area in Figure 2-(b)). See Appendix EC.3.6 for details.

3.3. Individual-level Implications of Interview-Stage Intervention

While the ρ-Rooney rule may change the group-level hiring outcomes (Theorem 1-(a)), its impact is not uniform across all candidates within the same group. Specifically, from Corollary 2, group m candidates with scores $a \in [a^*_m, a^\rho_m]$ lose interview opportunities, while group w candidates with scores $a \in [a^\rho_w, a^*_w]$ newly gain interview spots. In contrast, "strong" candidates with sufficiently high scores retain their interview spots under both (Benchmark) and the ρ-Rooney rule. Motivated by these heterogeneous changes in interview opportunities, we categorize candidates into groups based on their scores.

DEFINITION 1 **(Borderline vs. Strong Candidates).** *Given the ρ-Rooney rule, a candidate from group $i \in \{m, w\}$ is classified as follows:*

- ***Borderline candidate:*** *If $a \in [a^\rho_w, a^*_w]$ for group w and $a \in [a^*_m, a^\rho_m]$ for group m.*
- ***Strong candidate:*** *If $a > a^*_w$ for group w and $a > a^\rho_m$ for group m.*

[11] For function $f : \mathbb{R} \to \mathbb{R}$, its epigraph is defined as $\mathsf{epi}(f) := \{(x, r) \in \mathbb{R}^2 : f(x) \leq r\}$.

The ρ-Rooney rule has a straightforward effect on borderline candidates: it increases the hiring probability of borderline w candidates (from zero to positive) by granting them new interview opportunities, while decreasing the hiring probability of borderline m candidates (to zero) by reallocating their interview spots to the borderline w candidates. For strong candidates, the impact of the ρ-Rooney rule is more subtle. Although these candidates retain their interview opportunities, the rule changes the *composition* of the interview set by including borderline w candidates and excluding borderline m candidates. To understand how this shift affects strong candidates, it is helpful to first examine the hiring probability of candidates right at the interview threshold under (Benchmark).

Recall that the optimal interview thresholds must satisfy the balance condition (7). Under Assumption 1-(b) ($\tau_w > \tau_m$), this balance condition can only hold if $\bar{G}_m(s^*|a_m^*) > \bar{G}_w(s^*|a_w^*)$ (see Lemma EC.5-(a) in Section EC.3.9). In other words, at the optimal interview and hiring thresholds, group m candidates must have a higher hiring probability than their group w counterparts. Roughly speaking, this is because group w candidates at a_w^* are more likely to have extreme match values——either exceptionally high or low compared to their scores. In contrast, group m candidates' match values at a_m^* are more concentrated around their scores. Now, as the Rooney rule introduces more lower-score group w candidates into the interview set and excludes higher-score group m candidates, the composition of the interview set only exacerbates this disparity in the hiring probability. We formalize this result below.

LEMMA 1 **(Hiring Probability at Interview Thresholds)**. *For any $\rho \geq \rho^*$ where ρ^* is defined in (4), group m candidates at the interview threshold a_m^ρ have a higher hiring probability than group w candidates at a_w^ρ. Formally, $\bar{G}_m(s^\rho \mid a_m^\rho) > \bar{G}_w(s^\rho \mid a_w^\rho)$ for any $\rho \geq \rho^*$.*

We prove Lemma 1 in Appendix EC.3.10. Lemma 1 suggests that the adoption of the ρ-Rooney rule decreases candidate-level competition. To see this, suppose that the firm applies the ρ-Rooney rule with ρ being infinitesimally larger than ρ^*. This replaces group m candidates at threshold a_m^* (who had a higher hiring probability) with group w candidates at a_w^* (who had a lower hiring probability). As a result, the overall competition within the interview pool weakens. Formally, this shift lowers the hiring threshold, increasing the hiring probability for strong candidates in both groups. Furtheremore, as $\rho > \rho^*$ continues to increase, the firm's interview thresholds adjust continuously in ρ (Equation (8)), leading to a progressive weakening of candidate-level competition. In other words, the ρ-Rooney rule changes the composition of the interview set in a way that makes strong candidates appear more competitive relative to the new interview set, thereby increasing their hiring probabilities.

We formalize all the above observations in the following proposition, proven in Appendix EC.3.11.

PROPOSITION 3 **(Individual-level Impact of the Interview-Stage Intervention)**. *For any $\rho > \rho^*$ where ρ^* is defined in (4), the firm's adoption of the ρ-Rooney Rule results in the following:*

(a) ***Strong Candidates in Both Groups Gain:*** *The hiring probability increases for strong candidates in both groups. Specifically, $s^\rho \leq s^*$, with strict inequality whenever $s^* > 0$.*

(b) ***Borderline Group*** w ***Candidates Gain:*** *The hiring probability increases for borderline candidates in group* w.

(c) ***Borderline Group*** m ***Candidates Lose:*** *The hiring probability decreases for borderline candidates in group* m.

4. Two-Firm Labor Market

In this section, we extend our base model to a labor market with two vertically differentiated firms. We first introduce the extended model in Section 4.1 and characterize the unique equilibrium in Section 4.2. We then examine the impact of non-universal adoption of the interview-stage intervention on hiring outcomes at the market, both in group (Section 4.3) and individual (Section 4.4) levels.

4.1. Model

Extended Model Setup. A unit mass of candidates seeks employment at one of two firms, referred to as *Firm 1* and *Firm 2* who make interview and hiring decisions in parallel. The two firms are vertically differentiated in the sense that candidates have homogeneous preferences, strictly preferring Firm 1 over Firm 2. Candidates apply to both firms. The hiring process unfolds again in the two stages.

(i) Interviewing: Each Firm $f \in \{1,2\}$ selects which candidates to interview based on their scores and group identities. Formally, Firm f chooses an interview set $A_{f,i} \subseteq \mathbb{R}$, such that candidates from group i are interviewed by Firm f if $a \in A_{f,i}$. Each firm f has an interview capacity C_f. Notably, the two firm's interview sets may overlap since they compete on the same pool of candidates for hiring.

At the interview, Firm f observes the candidate's true match value, which is drawn from the conditional distribution $v \mid \{a,i\} \sim \mathcal{N}(a, \tau_i^2)$. We assume that, although the pre-interview scores are common across firms, conditional on a candidate's score a and group i, the match values are drawn independently for each firm. Similar to the single-firm model (Section 2), this assumption reflects that each firm observes its own idiosyncratic component of the candidate's match value during interviews.

(ii) Hiring: Each Firm $f \in \{1,2\}$ aims to maximize the expected match values of its hired candidates while respecting its hiring capacity Δ_f. Thus, each firm adopts a threshold-based hiring strategy, sending offers to candidates whose match values exceed a hiring threshold s_f to fill its hiring capacity Δ_f. Note that a candidate may receive offers from both firms; in this case, the candidate always accepts the offer from (the more-preferred) Firm 1. Thus, from Firm 1's perspective, nothing changes from the single-firm model as it does not face any competition. However, Firm 2 will take into account the firm-level competition when it chooses which candidates to interview and hire.

Benchmark vs. Intervention. Similar to the one-firm model, we consider two settings referred to as (Benchmark) and (Intervention). Under (Benchmark), both firms are unconstrained. Under (Intervention), Firm 1 adopts the ρ-Rooney rule, while Firm 2 remains unconstrained. As highlighted in Section 1, this reflects a practice where interview-stage interventions are voluntarily adopted by leading firms without a market-wide mandate (see Footnote 3). Similar to the single-firm model (Section 2), we

use ρ^* to denote the Firm 1's optimal interview fraction of group w under (Benchmark) (see Equation (4)). Recall that by Corollary 2, the ρ-Rooney rule changes Firm 1's decision only if $\rho > \rho^*$.

Firms' Optimization Problems and Equilibrium Concept. The above setup can be viewed as a static game in which each firm simultaneously decides its interview set $A_f = (A_{f,m}, A_{f,w})$ and hiring threshold s_f. For Firm 1, without the ρ-Rooney Rule, the dominant strategy is to ignore Firm 2 and solve the optimization problem (Benchmark). Thus, to characterize a Nash equilibrium of this game, it suffices to determine the best response of Firm 2, given Firm 1's decision. Concretely, under (Benchmark), Firm 1's optimal strategy (A_1^*, s_1^*) is characterized by Proposition 1: its optimal interview set is greedy in scores within each group, i.e., $A_{1,i}^* = [a_{1,i}^*, \infty)$ where interview threshold $\boldsymbol{a}_1^* = (a_{1,m}^*, a_{1,w}^*)$ and hiring threshold s_1^* are jointly determined by the optimality conditions (5) and (6). For simplicity, we interchangeably refer to Firm 1's interview strategy as its interview threshold $\boldsymbol{a}_1^*$. However, Firm 2's strategy depends on the pool of candidates not hired by Firm 1. Specifically, given Firm 1's optimal strategy $(\boldsymbol{a}_1^*, s_1^*)$, the fraction of candidates from group i with score a available to Firm 2 is given by:

$$\Psi_i(a \mid \boldsymbol{a}_1^*, s_1^*) = \mathbb{1}[a < a_{1,i}^*] + \mathbb{1}[a \geq a_{1,i}^*] G_i(s_1^* \mid a), \tag{13}$$

where $\mathbb{1}[\cdot]$ denotes the indicator function. We refer to $\Psi_i(\cdot \mid \boldsymbol{a}_1^*, s_1^*)$ as the *availability function*, which naturally reflects Firm 2's perceived competition for candidates. Notably, $\Psi_i(\cdot \mid \boldsymbol{a}_1^*, s_1^*)$ exhibits a jump discontinuity at Firm 1's interview threshold $a = a_{1,i}^*$. Firm 2's best response (A_2^*, s_2^*) is then determined by solving the following optimization program:

$$\mathsf{OPT}_2(\boldsymbol{a}_1^*, s_1^*) \quad \max_{\substack{A_{2,m}, A_{2,w} \subseteq \mathbb{R}, \\ s_2 \geq 0}} \sum_{i \in \{m,w\}} 0.5 \int_{A_{2,i}} \int_{s_2}^{\infty} v g_i(v \mid a) h_i(a) \Psi_i(a \mid \boldsymbol{a}_1^*, s_1^*) \, \mathrm{d}v \, \mathrm{d}a, \tag{14}$$

$$\text{s.t.} \sum_{i \in \{m,w\}} 0.5 \int_{A_{2,i}} h_i(a) \, \mathrm{d}a = C_2, \tag{15}$$

$$\sum_{i \in \{m,w\}} 0.5 \int_{A_{2,i}} \bar{G}_i(s_2 \mid a) h_i(a) \Psi_i(a \mid \boldsymbol{a}_1^*, s_1^*) \, \mathrm{d}a \leq \Delta_2. \tag{16}$$

We refer to this optimization program as $\mathsf{OPT}_2(\boldsymbol{a}_1^*, s_1^*)$. The objective function (14) maximizes the total match value of hired candidates, accounting for their availability $\Psi_i(\cdot \mid \boldsymbol{a}_1^*, s_1^*)$. Constraint (15) ensures that Firm 2 adheres to its interview capacity C_2, while constraint (16) limits the hiring mass, accounting for the availability, to at most Δ_2. Notably, Firm 2's optimization problem is considerably more challenging than Firm 1's, due to the interplay between its endogenous hiring threshold s_2—which depends on its interview set A_2—and the availability function $\Psi_i(\cdot \mid \boldsymbol{a}_1^*, s_1^*)$.

In the next subsection, we show that $\mathsf{OPT}_2(\boldsymbol{a}_1^*, s_1^*)$ has a unique optimal solution, implying that the static game has a unique Nash equilibrium. A Nash equilibrium under (Intervention) is defined analogously, with Firm 1's unconstrained strategy $(\boldsymbol{a}_1^*, s_1^*)$ replaced by its optimal strategy under the ρ-Rooney rule, which is denoted by $(\boldsymbol{a}_1^\rho, s_1^\rho)$ and characterized in Corollary 2. Correspondingly, Firm 2's optimal strategy in this case, denoted by (A_2^ρ, s_2^ρ), solves $\mathsf{OPT}_2(\boldsymbol{a}_1^\rho, s_1^\rho)$.

4.2. Structural Result

In this section, we characterize the unique equilibrium of the labor market introduced in Section 4.1. As discussed earlier, under (Benchmark), it suffices to characterize Firm 2's optimal solution that solves $\mathsf{OPT}_2(\boldsymbol{a}_1^*, s_1^*)$, which we establish in the following.

PROPOSITION 4 **(Characterization of Firm 2's Optimal Strategy)**. *Given Firm 1's optimal strategy* $(\boldsymbol{a}_1^*, s_1^*)$*, Firm 2's optimal strategy* (A_2^*, s_2^*) *that solves* $\mathsf{OPT}_2(\boldsymbol{a}_1^*, s_1^*)$ *uniquely exists and is jointly determined as follows:*

(a) ***Non-greedy Optimal Interview Set:*** *Let* $F_{2,i}(a, s_2^*) := \mathbb{E}[(v - s_2^*)_+ \mid a, i]\Psi_i(a \mid \boldsymbol{a}_1^*, s_1)$*. Then, Firm 2's optimal interview set* $A_{2,i}^*$ *for group* i *is a superlevel set of* $F_{2,i}(a, s_2^*)$*, i.e.,*

$$A_{2,i}^* = \{a : F_{2,i}(a, s_2^*) \geq \theta\}, \tag{17}$$

where level θ *is chosen to satisfy the interview capacity constraint* (15)*. Furthermore,* $A_{2,i}^*$ *is a union of two disjoint intervals: there exists end points* $(b_{2,i}^*, c_{2,i}, d_{2,i}^*)$ *such that* $-\infty < b_{2,i}^* \leq a_{1,i}^* \leq c_{2,i}^* \leq d_{2,i}^* < \infty$ *and*

$$A_{2,i}^* = [b_{2,i}^*, a_{1,i}^*] \cup [c_{2,i}^*, d_{2,i}^*]. \tag{18}$$

(b) ***Greedy Optimal Hiring Threshold:*** *The optimal hiring threshold* s_2^* *satisfies:*

$$s_2^* = \min\left\{s_2 \geq 0 : \sum_{i \in \{m,w\}} 0.5 \int_{A_{2,i}^*} \bar{G}_i(s_2 \mid a) h_i(a) \Psi_i(a \mid \boldsymbol{a}_1^*, s_1^*)\, \mathrm{d}a \leq \Delta\right\}. \tag{19}$$

Proposition 4-(a) establishes that the structure of Firm 2's optimal interview strategy is more complex than Firm 1's. Specifically, Firm 2's optimal interview set is a union of two intervals: a greedy lower interval covering candidates not interviewed by Firm 1, and a non-greedy (i.e., $d_{2,i}^*$ is finite) upper interval including candidates for which Firm 2 competes with Firm 1. Despite this distinct structure of the interview set from Firm 1, part (b) shows that Firm 2's hiring threshold is still greedy, selecting the interviewees with the highest non-negative match values until its hiring capacity Δ_2 is filled.

To understand the structure of Firm 2's interview set, it is helpful to understand the trade-off that Firm 2 faces when interviewing a high-score candidate. On the one hand, candidates with higher scores tend to have high match values. On the other hand, these candidates are also more likely to be hired by Firm 1. This trade-off naturally divides the candidate pool into two segments. Less competitive candidates who are not interviewed by Firm 1 ($a < a_{1,i}^*$) are always available to Firm 2; these candidates correspond to the lower greedy interval of the interview set. For candidates interviewed by Firm 1 ($a > a_{1,i}^*$), Firm 2 must carefully balance a candidate's diminishing availability with their potential for high match values. Because candidates with "exceptionally high" scores are only available to Firm 2 with vanishingly small probability, Firm 2 may not find worthwhile to use up its limited interview slots for such hard-to-get candidates. This leads Firm 2 to adopt a strategic and non-greedy interview strategy in the upper interval, targeting medium-score candidates who are more attainable.

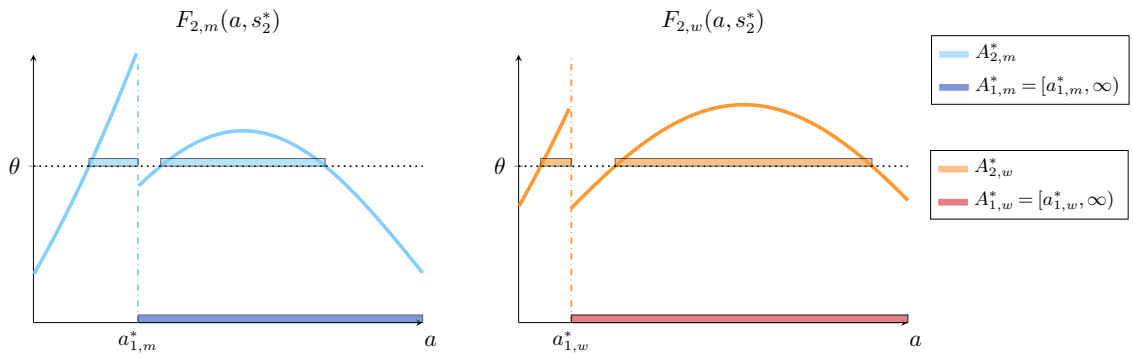


Figure 3 Illustration of Firm 2's optimal interview set (Proposition 4), which is a superlevel set of the discounted excess value function, $F_{2,i}(a, s_2^*) = \mathbb{E}\big[(v - s_2^*)_+ \mid a, i\big]\Psi_i(a \mid \boldsymbol{a}_1^*, s_1^*)$, with a common level θ across the two groups.

The above intuition is quantified through the *discounted excess value*, $F_{2,i}(a, s_2^*) := \mathbb{E}[(v - s_2^*)_+ \mid a, i]\Psi_i(a \mid \boldsymbol{a}_1^*, s_1^*)$, which discounts each candidate's excess value by their availability. Note that $F_{2,i}(a, s_2^*)$ depends on Firm 1's optimal strategy $(\boldsymbol{a}_1^*, s_1^*)$ through the availability $\Psi_i(a \mid \boldsymbol{a}_1^*, s_1^*)$, though we omit this dependence for brevity. Firm 2 greedily selects candidates based on their discounted excess value within its interview capacity. That is, the optimal interview set is a superlevel set of the discounted excess value with a common level θ across the two groups (Equation (17)). Therefore, characterizing $A_{2,i}^*$ boils down to understanding the behavior of the function $F_{2,i}(a, s_2^*)$ for different ranges of a. As illustrated in Figure 3, the function $F_{2,i}(a, s_2^*)$ is increasing for $a < a_{1,i}^*$. At $a = a_{1,i}^*$, there is a discontinuity in $F_{2,i}(a, s_2^*)$ due to the sudden transition in availability. For $a > a_{1,i}^*$, the shape of $F_{2,i}$ becomes highly nontrivial. Higher-score candidates have higher match values, but for the same reason, they are more likely to be hired by Firm 1. The crux of the proof for Proposition 4 is to show that $F_{2,i}$ is unimodal in score a, which implies the non-greedy structure of the superlevel set of $F_{2,i}(a, s^*)$. Roughly speaking, such unimodality arises because the excess value $\mathbb{E}[(v - s_2^*)_+ \mid a, i]$ increases linearly with score a, while the availability $\Psi(a \mid \boldsymbol{a}_1^*, s_1^*)$ decreases exponentially under normal distributions.

We prove Proposition 4 in Appendices EC.2 and EC.4.1. In Appendix EC.2, building on the same arguments used for Firm 1, we show that Equations (17) and (19) are necessary optimality conditions and that a solution satisfying these equations uniquely exists. In Appendix EC.4.1, we then establish that the discounted excess value function $F_{2,i}(a, s_2^*)$ is increasing in $a \le a_{1,i}^*$ and unimodal on $a \ge a_{1,i}^*$, which in turn implies the (non-greedy) structure of the interview set in Equation (18). Lastly, an analogous characterization can be established under (Intervention) by replacing Firm 1's strategy with $(\boldsymbol{a}_1^\rho, s_1^\rho)$. Specifically, Firm 2's optimal interviews set under (Intervention) is given by $A_{2,i}^\rho = [b_{2,i}^\rho, a_{1,i}^\rho] \cup [c_{2,i}^\rho, d_{2,i}^\rho]$ where $-\infty < b_{2,i}^\rho \le a_{1,i}^\rho \le c_{2,i}^\rho \le d_{2,i}^\rho < \infty$.

Remark 1 (Practical Appeal). The non-greedy structure of Firm 2's interview strategy aligns well with interview patterns observed in practice. In markets with simultaneous interviews, less competitive firms may forgo interviewing outstanding candidates. The academic job market is a representative example: schools conduct interviews in parallel from the same pool of candidates. In Appendix EC.6,

we use data from the academic job market in economics (EJMR 2023) to illustrate lower-tier schools tend to interview fewer outstanding candidates interviewed by multiple top-tier schools.

Remark 2 (Beyond Normal Distributions). The non-greedy interview also emerges under other conditional match value distributions. Notably, in Appendix EC.5.2, we show that if conditional match values follow a Gumbel distribution (a common noise structure in consumer choice modeling), the lower-ranked firm's optimal interview strategy remains non-greedy, similar to Proposition 4.

4.3. Group-level Implications of Non-universal Interview-Stage Intervention

In this section, we study the impact of Firm 1's adoption of the ρ-Rooney Rule on group-level hiring outcomes in the labor market. Having already analyzed the intervention's impact on Firm 1 itself in Section 3, we focus on its *downstream* impact on Firm 2. Specifically, we ask: How does Firm 2's optimal strategy change from (Benchmark) to (Intervention)?

To build intuition, consider the special case where $C_f = \Delta_f$. In this regime, both firms' optimal hiring thresholds must be zero regardless of the ρ-Rooney rule. This is because, for any interview set with mass $C_f = \Delta_f$, the mass of interviewees with non-negative match values is strictly less than Δ_f. As a result, both firms must set the lowest possible hiring threshold (i.e., zero) to satisfy the optimality condition of the hiring threshold (Equations (6) and (19)). This implies that firm-level competition is at its *strongest* for Firm 2, as most candidates interviewed by Firm 1 are unavailable to Firm 2. This extreme scarcity makes Firm 2 pessimistic about hiring candidates from Firm 1's interview set. Formally, the discounted excess value $F_{2,i}(a,0)$ for these candidates is so small that its superlevel set, characterizing Firm 2's optimal interview set (see Proposition 4), excludes them entirely (see Figure 4). (For $C_f = \Delta_f$, function $F_{2,i}(a,0)$ remains unimodal for $a > a^*_{1,i}$. However, its maximum now occurs before $a^*_{1,i}$, and thus the function illustrated in Figure 4 is decreasing in $a > a^*_{1,i}$.)

This special regime provides a clear lens to understand the downstream impact. We illustrate in Figure 4 how Firm 2's optimal interview set changes from (Benchmark) to (Intervention). From Proposition 4, the optimal interview set in either setting is a superlevel set of the discounted excess value function $F_{2,i}(a,0)$. When Firm 1 adopts the ρ-Rooney rule, its interview threshold shifts from $\boldsymbol{a}^*_1$ to $\boldsymbol{a}^\rho_1$ such that $a^*_{1,m} < a^\rho_{1,m}$ and $a^*_{1,w} > a^\rho_{1,w}$ (see Corollary 2). This directly affects the availability function $\Psi_i(a \mid \boldsymbol{a}^*_1, 0)$ and, consequently, the discounted excess value function $F_{2,i}(a,0)$.

Specifically, using categorization of Definition 1, the ρ-Rooney rule replaces borderline m candidates in Firm 1's interview set with borderline w candidates. From Firm 2's perspective, this replacement significantly increases the availability of borderline m candidates, raising their discounted excess value. For illustration, see Figure 4 that compares the group m's discounted excess value under (Benchmark) (thin, light blue) and (Intervention) (thick, dark blue). Conversely, borderline w candidates are now included in Firm 1's interview set, thus similarly decreasing their discounted excess value for Firm 2. See the right panel of Figure 4 that compares the group w's discounted excess value under (Benchmark)

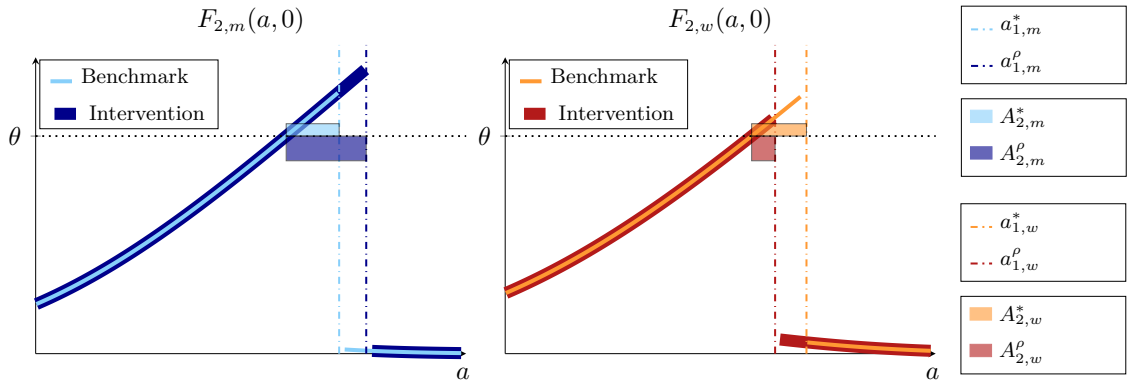


Figure 4 **Change in Firm 2's interview strategy due to Firm 1's adoption of the ρ-Rooney rule when $C_f = \Delta_f$.**

(thin, light orange) and (Intervention) (thick, dark red). Consequently, Firm 2 replaces *all* borderline w candidates with borderline m candidates in its interview set. As a result, Firm 2 interviews (and therefore hires) strictly fewer group w candidates. We formalize this intuition in the following.

PROPOSITION 5 **(Downstream Impact of Firm 1's Adoption of the ρ-Rooney Rule).** *For small enough $C_f = \Delta_f$, $f \in \{1,2\}$, Firm 1's adoption of the ρ-Rooney rule for all $\rho > \rho^*$ strictly decreases Firm 2's hiring mass of group w. Formally, let $\lambda^*_{2,w}$ and $\lambda^\rho_{2,w}$ denote Firm 2's hiring mass of group w under (Benchmark) and (Intervention), respectively. Then, $\lambda^\rho_{2,w} < \lambda^*_{2,w}$ for any $\rho > \rho^*$.*

We prove Proposition 5 in Appendix EC.4.2. Recall from Equation (17) that Firm 2's optimal interview set is the superlevel set of the function $F_{2,i}(a,0)$. Let θ be the corresponding level in Equation (17). In the proof, we show that, this level θ *does not change* from (Benchmark) to (Intervention) under the special regime $C_f = \Delta_f$. This invariance in θ translates to the "exchange" of group w borderline candidates with group m in Firm 2's interview set, as in Figure 4.

Obtaining an analytical characterization of the ρ-Rooney rule's downward impact in a more general case (i.e., when $C_f > \Delta_f$ for either firm $f \in \{1,2\}$) is significantly more challenging due to the endogenous nature of the optimal hiring threshold s^*_2. However, our extensive numerical simulations indicate that our insights extend beyond this special case. As an illustrative example, in Figure 5, we present the change in the hiring fraction of group w within Firm 2, as well as in the overall market, when moving from (Benchmark) to (Intervention) with $\rho = 0.5$. For this simulation, we assume that the two firms have identical interview capacity C and hiring capacity Δ, and re-use the same parameters used in Figure 2. By Theorem 1, the 0.5-Rooney rule has a non-zero impact only when C is sufficiently small. The 0.5-Rooney rule then forces Firm 1 to adjust its strategy, prompting Firm 2 to optimally respond by decreasing its hires from group w (see the left panel of Figure 5). More strikingly, this downward impact outweighs the diversity improvement within Firm 1, leading to a net reduction in group w's representation in the overall market, as seen in the right panel of Figure 5. This sharply contrasts our finding in the single-firm model: despite being well-intended to improve representation within the top firm, the non-universal adoption of the Rooney Rule ultimately reduces the group w's

representation in the market due to firm-level competition. In the next subsection, we explain the underlying mechanism behind this market-wide decline in group w's representation.

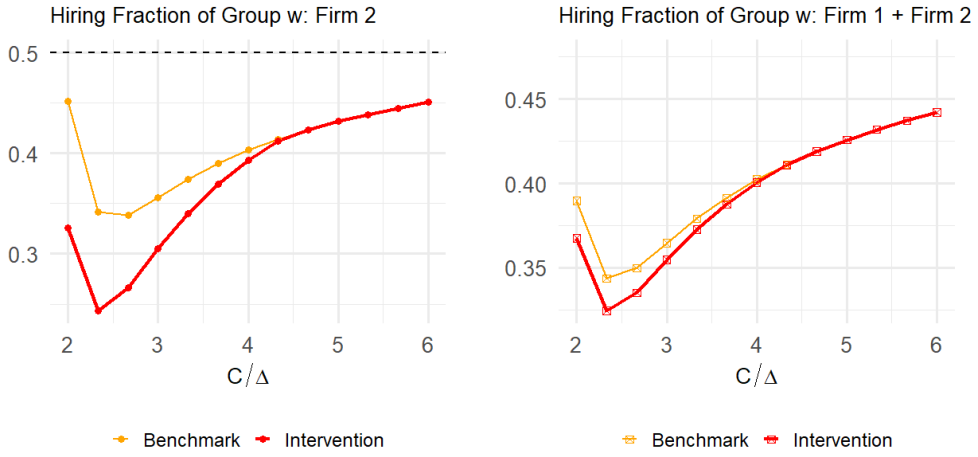


Figure 5 **Group w's hiring fraction in Firm 2 (left panel) and in the overall market (right panel) with $\rho = 0.5$. Both firms have identical capacity parameters C and Δ. We vary C and reuse the other parameters from Figure 2-(a).**

4.4. Individual-level Implications of Non-universal Interview-Stage Intervention

We now turn to the individual-level implications of Firm 1's adoption of the ρ-Rooney rule. Similar to the single-firm model (Section 3.3), the intervention has a different impact on different groups of candidates based on their scores. Note that a candidate's hiring probability in the market depends on which firm they are interviewed by. In light of this, we introduce a refined category of candidates called *superstars*, whose interview spot remains the same regardless of the intervention. These are candidates with exceptionally high scores who are interviewed exclusively by Firm 1, as Firm 2 bypasses them due to the low probability of attaining them. We formally define superstar candidates below.

Definition 2 **(Superstar Candidate)**. *A superstar candidate from group i with score a is a strong candidate (see Definition 1) who is interviewed exclusively by Firm 1 under both* (Benchmark) *and* (Intervention)*. Formally, a group i candidate is classified as a superstar if their score a satisfies $a \geq \max\{d^*_{2,i}, d^\rho_{2,i}\}$ where $d^*_{2,i} = \max A^*_{2,i}$ and $d^\rho_{2,i} = \max A^\rho_{2,i}$.*

Note that the superstar candidates always exist since, from Proposition 4, $\max\{d^*_{2,i}, d^\rho_{2,i}\}$ is finite. As superstars are only interviewed by Firm 1, their hiring probability in the market coincides with their hiring probability by Firm 1. Consequently, the implications of (Intervention) on the hiring probability for superstar candidates follow naturally from the single-firm model (Proposition 3). That is, the ρ-Rooney rule weakens *candidate-level* competition within Firm 1's interview set ($s^*_1 \geq s^\rho_1$), making superstar candidates appear more competitive and increasing their hiring probabilities.

We now turn to the borderline candidates (see Definition 1). Unlike superstars, the borderline candidates may experience shifts in their interview spots within the market. For example, in the most

competitive setting of $C_f = \Delta_f$ (as we discussed in Section 4.3), *all* borderline w candidates lose their spots with Firm 2 but gain spots with Firm 1, while *all* borderline m candidates experience the opposite shift. In more general regimes, we show that similar changes occur for a subset of borderline candidates.

PROPOSITION 6 **(Interview Advantage & Disadvantage for Borderline Candidates).** *For small enough Δ_f, $f \in \{1,2\}$, there exist interview capacity thresholds $\bar{C}_f$ such that if $C_f \in [\Delta_f, \bar{C}_f]$, Firm 1's adoption of the ρ-Rooney Rule for all $\rho > \rho^*$ results in the following:*

1. ***Interview Advantage for Borderline** w **Candidates:** a positive mass of borderline group w candidates loses an interview spot from Firm 2 but gains one from Firm 1.*
2. ***Interview Disadvantage for Borderline** m **Candidates:** a positive mass of borderline group m candidates loses an interview spot from Firm 1 but gains one from Firm 2.*

We prove Proposition 6 in Appendix EC.4.3. At first glance, this change of interview spots may seem beneficial for borderline group w candidates. However, this advantage in the interview stage may not lead to a better hiring outcome because the intensity of candidate-level competition differs across firms. Intuitively, Firm 1 may be more selective in hiring as it greedily interviews the high-score candidates, leading to stronger candidate-level competition compared to Firm 2. In this case, borderline group w candidates gaining an interview spot at Firm 1 may face a *decrease* in hiring probability due to fiercer candidate-level competition, while borderline group m candidates gaining an interview spot at Firm 2 may see an *increase* due to weaker candidate-level competition.

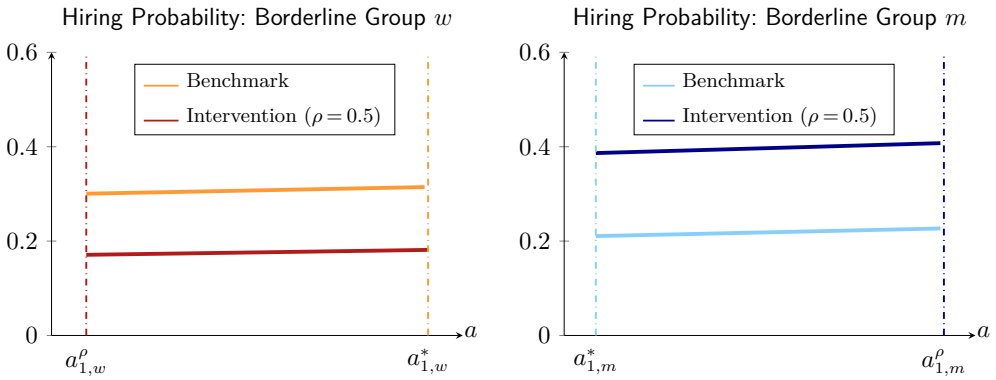


Figure 6 Changes in the hiring probability for borderline candidates. Both firms have identical capacity parameters $C = 0.03$ and $\Delta = 0.01$. We reuse the other parameters from Figure 5.

Figure 6 confirms this intuition through simulation. Here, we use the same parameters as in Figure 5. Focusing on $C/\Delta = 3$, we plot the hiring probability of all borderline group m candidates (resp., group w) in the left (resp., right) panel. In this simulation, all borderline group w candidates lose an interview spot at Firm 2 and gain one at Firm 1 under (Intervention). However, they face fiercer candidate-level competition due to Firm 1's higher hiring threshold, $s^{\rho}_1 > \max\{s^*_2, s^{\rho}_2\}$ (note $\min\{s^{\rho}_1, s^*_1\} = s^{\rho}_1$ by

Proposition 3), leading to a uniform decrease in their hiring probabilities. Conversely, borderline group m candidates gain an interview spot at Firm 2 and experience an increase in their hiring probabilities.

We formalize these observations in the following proposition, which we prove in Appendix EC.4.4:

PROPOSITION 7 **(Individual-level Implication of Firm 1's Adoption of the ρ-Rooney Rule).** *For any $\rho > \rho^*$, there exist Firm 2's interview capacity thresholds $\bar{C}_2$ such that if $C_2 \in [\Delta_2, \bar{C}_2]$, Firm 1's adoption of the ρ-Rooney Rule results in the following:*

(a) ***Superstar Candidates in Both Groups Gain:*** *The hiring probability increases for superstar candidates in both groups ($s_1^* \geq s_1^\rho$).*

(b) ***Some Borderline Group w Candidates Lose:*** *For any borderline group w candidates who lose an interview spot from Firm 2, the hiring probability decreases ($s_1^\rho \geq s_2^*$).*

(c) ***Some Borderline Group m Candidates Gain:*** *For any borderline group m candidates who gain an interview spot from Firm 2, the hiring probability increases ($s_1^* \geq s_2^\rho$).*

To further complement our theoretical analysis, we use simulations across a wide range of (C_1, C_2) to examine the extent to which borderline candidates are impacted under the intervention. In Figure 7, we examine two subsets of borderline candidates: D_w, the set of borderline group w candidates who lose an interview spot at Firm 2 but gain one at Firm 1, and D_m, the set of borderline group m candidates who lose an interview spot at Firm 1 but gain one at Firm 2. The size of each circle represents the mass of D_w (left panel) and D_m (right panel), normalized by the mass of borderline candidates in each group. The intensity of each circle's color reflects the magnitude of the guaranteed change in hiring probabilities: the maximum change for D_w and the minimum change for D_m. For example, if the maximum change for D_w is $-\epsilon$ for $\epsilon > 0$, then all candidates in D_w experience a decrease in hiring probability of at least ϵ. In the left panel, red indicates a decrease in hiring probability for *all* candidates in D_w, while in the right panel, blue indicates an increase for D_m. Figure 7 confirms that significant proportion of borderline group w (resp. group m) candidates simultaneously experience shifts in the interview spots and uniform decrease (resp. increase) in hiring probabilities.

We conclude this section with two remarks. First, note that the implications for borderline group w candidates under the two-firm model contrast with those in the single-firm model: they uniformly benefit in the single-firm model (Proposition 3), but may be harmed in the two-firm model. This "reversal" highlights the importance of a market-wide view in studying the interview-stage interventions. From Firm 1's perspective, the intervention appears to favor borderline w candidates by *newly* offering them interview spots. This was beneficial in the single-firm model, where they otherwise had *no chance* of being hired. However, in the two-firm model, these candidates become vulnerable to Firm 2's strategic response: they are "competitive enough" to secure spots at Firm 2 under the benchmark but lose them when Firm 1 begins interviewing them. As a result, the intervention simply shifts their interviews from Firm 2 to Firm 1, where they face tougher candidate-level competition from higher-scoring candidates.

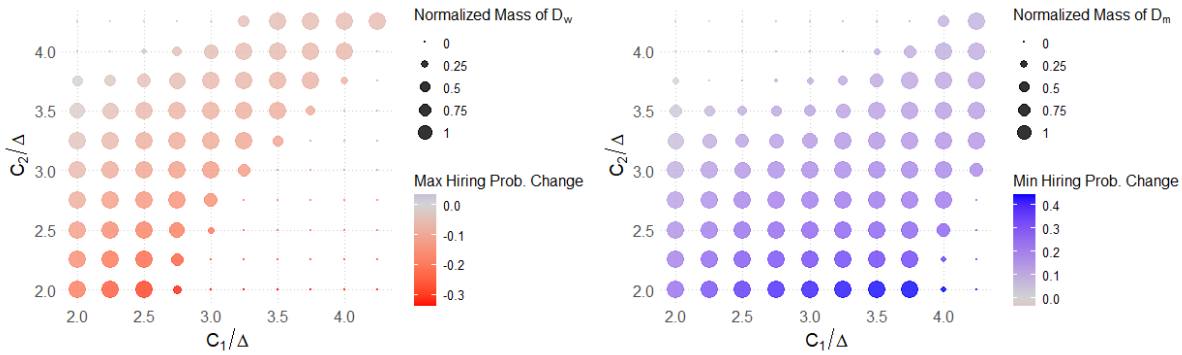


Figure 7 Changes of the interview spots (Proposition 6) and the resulting changes in hiring probabilities (Proposition 7) for borderline candidates from (Benchmark) to (Intervention) with $\rho = 0.5$. We reuse the parameters from Figure 5 with fixed $\Delta = 0.01$, but vary (C_1, C_2).

Thus, while Firm 1 intended to benefit borderline w candidates, the intervention ultimately harmed them due to Firm 2's strategic response. Lastly, this mechanism is a key driver of the overall decrease in group w's hiring rate, as illustrated in Figure 5: due to stronger candidate-level competition in Firm 1, the increase in group w's hiring rate at Firm 1 may not fully offset the decrease at Firm 2.

5. Conclusion

Motivated by the study of interview-stage interventions, such as the Rooney rule, we developed a novel model of the labor market featuring a two-stage hiring process and pre-interview information asymmetry across different social groups, explicitly incorporating capacity constraints for both hiring and interviews. For a single firm, we showed that even in an idealized setting where, post-interview, no asymmetry exists between the two groups, disparity in hiring outcomes arises due to limited interview capacity. Our findings align with empirical evidence documenting the limited impact of these policies (Cecchi-Dimeglio 2022, Smith and Lasker 2023) and shed new light on the challenges of improving diversity in hiring outcomes through soft affirmative actions. Another important insight is that strong candidates from both groups emerge as beneficiaries of such interventions, as they now face weaker candidate-level competition. In light of our results, a natural future direction is to explore alternative diversity policies beyond interview-stage interventions and examine how their impact interacts with realistic operational considerations such as capacity constraints.

Moving beyond the single-firm case, we investigate the impact of firm-level competition in a vertically differentiated market where only the top firm adopts the intervention. We show that such non-universal adoption, commonly practiced across different industries, can have adverse effects on hiring outcomes: the policy may reduce group-level diversity at the lower-ranked firm by incentivizing it to hire fewer targeted candidates due to increased firm-level competition. Furthermore, some candidates from the very group the policy intends to help may ultimately be harmed, as they face fiercer candidate-level competition at the top firm. Overall, our findings suggest that, in addition to limited interview capacity,

firm-level competition is another key operational consideration that limits the effectiveness of interview-stage interventions. Extending our framework to labor markets with different characteristics, such as horizontally differentiated firms or settings with non-simultaneous hiring decisions, would be an interesting avenue for future research.

References

Ahmed Agiza, Mohamed Mostagir, and Sherief Reda. Analyzing the impact of data selection and fine-tuning on economic and political biases in llms. *arXiv e-prints*, pages arXiv–2404, 2024.

Charalambos D Aliprantis and Owen Burkinshaw. *Principles of real analysis*. Gulf Professional Publishing, 1998.

Amine Allouah, Christian Kroer, Xuan Zhang, Vashist Avadhanula, Nona Bohanon, Anil Dania, Caner Gocmen, Sergey Pupyrev, Parikshit Shah, Nicolas Stier-Moses, et al. Fair allocation over time, with applications to content moderation. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 25–35, 2023.

Mohammad Reza Aminian, Vahideh Manshadi, and Rad Niazadeh. Markovian search with socially aware constraints. *Available at SSRN 4347447*, 2023.

Kenneth Arrow. The Theory of Discrimination. 1971.

Itai Ashlagi, Jiale Chen, Mohammad Roghani, and Amin Saberi. Stable matching with interviews. *arXiv preprint arXiv:2501.12503*, 2025.

Jackie Baek and Ali Makhdoumi. The feedback loop of statistical discrimination. *Available at SSRN 4658797*, 2023.

Ariane Baye and Christian Monseur. Gender differences in variability and extreme scores in an international context. *Large-scale Assessments in Education*, 4:1–16, 2016.

Gary S Becker. The Economics of Discrimination. *University of Chicago Press*, 1957.

Moshe Ben-Akiva. Discrete choice analysis: Theory and application to travel demand. *MIT Press google schola*, 2:575–589, 1985.

Dimitris Bertsimas, Vivek F Farias, and Nikolaos Trichakis. The price of fairness. *Operations research*, 59(1): 17–31, 2011.

Patrick Billingsley. *Probability and Measure*. Wiley Series in Probability and Mathematical Statistics. Wiley, New York, 3rd edition, 1995.

Carlos Bonet, Nick Arnosti, and Jay Sethuraman. Explainable affirmative action. In *Proceedings of the 25th ACM Conference on Economics and Computation*, pages 310–310, 2024.

Paola Cecchi-Dimeglio. Why the 30 Percent Mansfield Rule Can't Work: A Supply-Demand Empirical Analysis of Leadership in the Legal Profession. *Fordham Law Review*, 91:1161, 2022.

L Elisa Celis, Chris Hays, Anay Mehrotra, and Nisheeth K Vishnoi. The Effect of the Rooney Rule on Implicit Bias in the Long Term. In *Proceedings of the 2021 ACM conference on fairness, accountability, and transparency*, pages 678–689, 2021.

Hongqiao Chen, Ming Hu, and Jing Peng. Socially responsible pricing. *Available at SSRN 4490937*, 2023.

Qinyi Chen, Negin Golrezaei, and Fransisca Susan. Fair assortment planning. *arXiv preprint arXiv:2208.07341*, 2022.

Maxime C Cohen, Adam N Elmachtoub, and Xiao Lei. Price discrimination with fairness constraints. *Management Science*, 68(12):8536–8552, 2022.

Bo Cowgill. Bias and productivity in humans and algorithms: Theory and evidence from resume screening. *Columbia Business School, Columbia University*, 29, 2020.

Diversity Lab. *Commitment to Widening the Leadership Talent Pool Continues to Grow as 240+ Law Firms Become Mansfield Certified*, 2023. URL `https://www.diversitylab.com/wp-content/uploads/2023/10/2022-2023-U.S.-Canada-Mansfield-Certified-Large-Midsize-Firm-Press-Release-October-2023.pdf`. Accessed: 2025-03-03.

Federico Echenique, Ruy Gonzalez, Alistair J Wilson, and Leeat Yariv. Top of the batch: Interviews and the match. *American Economic Review: Insights*, 4(2):223–238, 2022.

Florian Ederer, Paul Goldsmith-Pinkham, and Kyle Jensen. Anonymity and identity online. *arXiv preprint arXiv:2409.15948*, 2024.

EJMR. Economics, math and sociology job market wiki, 2023. URL `https://www.econjobrumors.com/wiki2023.php`.

Vitalii Emelianov, Nicolas Gast, Krishna P Gummadi, and Patrick Loiseau. On fair selection in the presence of implicit variance. In *Proceedings of the 21st ACM Conference on Economics and Computation*, pages 649–675, 2020.

Alessandro Fabris, Nina Baranowska, Matthew J Dennis, David Graus, Philipp Hacker, Jorge Saldivar, Frederik Zuiderveen Borgesius, and Asia J Biega. Fairness and bias in algorithmic hiring: A multidisciplinary survey. *ACM Transactions on Intelligent Systems and Technology*, 2024.

Hanming Fang and Andrea Moro. Theories of statistical discrimination and affirmative action: A survey. *Handbook of Social Economics*, 1:133–200, 2011.

Daniel Fershtman and Alessandro Pavan. "soft" affirmative action and minority recruitment. *American Economic Review: Insights*, 3(1):1–18, 2021.

Bailey Flanigan, Paul Gölz, Anupam Gupta, Brett Hennig, and Ariel D Procaccia. Fair algorithms for selecting citizens' assemblies. *Nature*, 596(7873):548–552, 2021.

Rupert Freeman, Seyed Majid Zahedi, and Vincent Conitzer. Fair and efficient social choice in dynamic settings. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence*, pages 4580–4587, 2017.

Sarah Frier. Does Rooney Rule Work? C-Suite Hiring Isn't More Diverse. *Bloomberg*, 2022.

Nikhil Garg, Hannah Li, and Faidra Monachou. Dropping standardized testing for admissions trades off information and access. *arXiv preprint arXiv:2010.04396*, 2020.

Nikhil Garg, Wes Gurnee, David Rothschild, and David Shmoys. Combatting gerrymandering with social choice: The design of multi-member districts. In *Proceedings of the 23rd ACM Conference on Economics and Computation*, pages 560–561, 2022.

Martin Gonzalez-Cabello, Auyon Siddiq, Charles Corbett, and Catherine Hu. Fairness in crowdwork: Making the human ai supply chain more humane. *Business Horizons*, 09 2024. doi: 10.1016/j.bushor.2024.09.003.

Swati Gupta and Vijay Kamble. Individual fairness in hindsight. *Journal of Machine Learning Research*, 22 (144):1–35, 2021.

Jonathan Guryan and Kerwin Kofi Charles. Taste-based or statistical discrimination: The economics of discrimination returns to its roots. *The Economic Journal*, 123(572):F417–F432, 2013.

Lily Hu and Yiling Chen. A short-term intervention for long-term fairness in the labor market. In *Proceedings of the 2018 World Wide Web Conference*, pages 1389–1398, 2018.

Ming Hu, Yanzhi Li, and Tongwen Wu. Diversity-fair online selection. *arXiv preprint arXiv:2504.10389*, 2025.

Sangram V Kadam. Interviewing in matching markets. 2015.

Sampath Kannan, Aaron Roth, and Juba Ziani. Downstream effects of affirmative action. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*, pages 240–248, 2019.

Eunjee Kim, Xiaoyang Long, and Jordan Tong. Fair funnels: Bias, performance, and interventions in multistage hiring processes. *Performance, and Interventions in Multistage Hiring Processes (January 15, 2025)*, 2025.

Jon Kleinberg and Manish Raghavan. Selection problems in the presence of implicit bias. *arXiv preprint arXiv:1801.03533*, 2018.

Junpei Komiyama and Shunya Noda. On statistical discrimination as a failure of social learning: A multiarmed bandit approach. *Management Science*, 2024.

Mimi Lamarre. More Than 360 Law Firms Gain Mansfield Certification, With Firms Remaining Quietly Committed to Diversity. *The American Lawyer*, 2024.

Tomás Larroucau, Ignacio Rios, Anaïs Fabre, and Christopher Neilson. College application mistakes and the design of information policies at scale. *Unpublished paper, Arizona State University, Tempe*, 2024.

Robin S Lee and Michael Schwarz. Interviewing in two-sided matching markets. *The RAND Journal of Economics*, 48(3):835–855, 2017.

David Li and Diana Dasrath. Former Dolphins Coach Brian Flores Files Racial Discrimination Lawsuit Against NFL, New York Giants. *NBC News*, 2022.

Stephen Machin and Tuomas Pekkarinen. Global sex differences in test score variability. *Science*, 322(5906): 1331–1332, 2008.

Vikram Manjunath and Thayer Morrill. Interview hoarding. *Theoretical Economics*, 18(2):503–527, 2023.

Vahideh Manshadi, Rad Niazadeh, and Scott Rodilitz. Fair dynamic rationing. *Management Science*, 69(11): 6818–6836, 2023a.

Vahideh Manshadi, Scott Rodilitz, Daniela Saban, and Akshaya Suresh. Redesigning volunteermatch's ranking algorithm: Toward more equitable access to volunteers. *Available at SSRN 4497747*, 2023b.

Andreu Mas-Colell, Michael Whinston, and Jerry Green. *Microeconomic Theory*. Oxford University Press, 1995. URL `https://EconPapers.repec.org/RePEc:oxp:obooks:9780195102680`.

Faidra Monachou and Itai Ashlagi. Discrimination in online markets: Effects of social bias on learning from reviews and policy design. *Advances in Neural Information Processing Systems*, 32, 2019.

Alfred Müller. Stochastic ordering of multivariate normal distributions. *Annals of the Institute of Statistical Mathematics*, 53:567–575, 2001.

NFL Football Operations. The Rooney Rule. `https://operations.nfl.com/inside-football-ops/inclusion/the-rooney-rule/`, 2025.

Paula Onuchic. Recent contributions to theories of discrimination. *arXiv preprint arXiv:2205.05994*, 2022.

Neil Paine. NFL Teams Are Making Brian Flores's Case for Him. *FiveThirtyEight*, Feb 2022. URL `https://fivethirtyeight.com/features/nfl-teams-are-making-brian-floress-case-for-him/`.

Edmund S Phelps. The statistical theory of racism and sexism. *American Economic Review*, 62(4):659–661, 1972.

Lincoln Quillian, Devah Pager, Ole Hexel, and Arnfinn H Midtbøen. Meta-analysis of field experiments shows no change in racial discrimination in hiring over time. *Proceedings of the National Academy of Sciences*, 114(41):10870–10875, 2017.

Christopher I Rider, James B Wade, Anand Swaminathan, and Andreas Schwab. Racial disparity in leadership: Evidence of valuative bias in the promotions of National Football League coaches. *American Journal of Sociology*, 129(1):227–275, 2023.

Jesse M Rothstein. College performance predictions and the SAT. *Journal of Econometrics*, 121(1-2):297–317, 2004.

Jad Salem and Swati Gupta. Secretary problems with biased evaluations using partial ordinal information. *Management Science*, 70(8):5337–5366, 2024.

Michael R Sampford. Some inequalities on mill's ratio and related functions. *The Annals of Mathematical Statistics*, 24(1):130–132, 1953.

Moshe Shaked and J George Shanthikumar. *Stochastic orders*. Springer, 2007.

Allen Sirolly, Yash Kanoria, and Hongyao Ma. The impact of race-blind and test-optional admissions on racial diversity and merit. *Available at SSRN 4873387*, 2024.

DeMaurice F Smith and Carl Lasker. The Rooney Suggestion: How the "Rule" Has Failed to Defeat Institutional Barriers to Equitable Hiring Practices in the NFL and Recommendations for Meaningful Reform. *Yale Law and Poliy Review*, 42:238, 2023.

Spencer Stuart. 2023 New Director Diversity Snapshot. `https://www.spencerstuart.com/-/media/2023/august/usbi2023/director_diversity/2023_new_director_-diversity_snapshot.pdf`, 2023.

Nancy L Stokey and Robert E Lucas Jr. *Recursive methods in economic dynamics*. Harvard University Press, 1989.

U.S. Census Bureau. Offices of lawyers (naics 54111), 2025. URL `https://data.census.gov/profile/54111_-_Offices_of_lawyers`. Accessed: 2025-03-03.

U.S. News & World Report. Best economics schools rankings, 2024. URL `https://www.usnews.com/best-graduate-schools/top-humanities-schools/economics-rankings`. Accessed: March 4, 2025.

Akhil Vohra and Nathan Yoder. Matching with costly interviews: The benefits of asynchronous offers. 2023.

Caleb Wong. UT System Implements "Rooney Rule," Promotes Diversity in Faculty, Staff. *The Daily Texan*, 2016.

EC.1. Table of Notation

The following table lists major notation used throughout the main body. For a single-firm model (Section 2), we omit the subscript f for brevity.

Table EC.1 Mathematical Notation

Symbol	Meaning
$C_f \in (0,1)$	interview mass of Firm f
$\Delta_f \in (0,1)$	hiring mass of Firm f
$i \in \{m,w\}$	social group
a	(pre-interview) score, $a \sim \mathcal{N}(0,\sigma_i^2)$
σ_i^2	variance of group i's score
$H_i(a)$ (resp., $h_i(a)$)	C.D.F. (resp., p.d.f.) of group i's pre-interview score
Φ (resp., ϕ)	C.D.F. (resp., p.d.f.) of standard Normal distribution
$(x)_+$	$\max(x,0)$
v	match value, $v \mid \{a,i\} \sim \mathcal{N}(a,\tau_i^2)$
τ_i^2	variance of group i's conditional match value distribution
$G_i(v \mid a)$ (resp., $g_i(v \mid a)$)	C.D.F. (resp., p.d.f.) of group i's conditional match value
$\Psi_i(a \mid \boldsymbol{a}_1^*, s_1^*)$	Availability of group i candidates with score a to Firm 2, given Firm 1's optimal strategy $(\boldsymbol{a}_1^*, s_1^*)$ (see Equation (13))
$A_f = (A_{f,m}, A_{f,w})$	Firm f's interview set for each group $i \in \{m,w\}$
s_f	Firm f's hiring threshold
(A_f^*, s_f^*)	Firm f's equilibrium strategy under (Benchmark)
ρ	minimum fraction of group w in the interview set required by the ρ-Rooney rule (see (ρ-Rooney rule))
(A_f^ρ, s_f^ρ)	Firm f's equilibrium strategy under Firm 1's adoption of the ρ-Rooney rule
$a_{1,i}^*$ (resp., $a_{1,i}^\rho$)	Firm 1's optimal interview threshold under (Benchmark) (resp., the ρ-Rooney rule) — see Proposition 1 and Corollary 2
ρ^*	Firm 1's interview fraction of group i under (Benchmark) (Equation (4))
$(b_{2,i}^*, c_{2,i}^*, d_{2,i}^*)$	end points of Firm 2's interview set $A_{2,i}^*$ for group i under (Benchmark) (i.e., $A_{2,i}^* = [b_{2,i}^*, a_{1,i}^*] \cup [c_{2,i}^*, d_{2,i}^*]$; see Proposition 4)
$(b_{2,i}^\rho, c_{2,i}^\rho, d_{2,i}^\rho)$	end points of $A_{2,i}^\rho$
$\lambda_{2,w}^*$ (resp., $\lambda_{2,w}^\rho$)	Firm 2's hiring mass of group w under (Benchmark) (resp., the ρ-Rooney rule)

EC.2. Meta-Characterization of Equilibrium Structure

In this appendix, we present meta-applicable structural results for the following general model. We will frequently use this meta-characterization to derive the structural results for the single-firm model (Proposition 1; Appendix EC.3.2) and the two-firm model (Proposition 4; Appendix EC.4.1), as well as additional structural results beyond our context (Appendices EC.5.1 and EC.5.2).

Specifically, consider the following optimization problem (OPT-Meta):

$$\text{(OPT-Meta)} \quad \max_{\substack{B_m, B_w \subseteq \mathbb{R} \\ s \geq 0}} \sum_{i \in \{m,w\}} 0.5 \int_{B_i} \int_s^\infty v g_i(v \mid a) h_i(a) \Psi_i(a) \, \mathrm{d}v \, \mathrm{d}a \tag{EC.1}$$

$$\text{s.t.} \sum_{i \in \{m,w\}} 0.5 \int_{B_i} h_i(a) \, \mathrm{d}a = C \tag{EC.2}$$

$$\sum_{i \in \{m,w\}} 0.5 \int_{B_i} \int_s^\infty g_i(v \mid a) h_i(a) \Psi_i(a) \, \mathrm{d}v \, \mathrm{d}a \leq \Delta \tag{EC.3}$$

To see the connection between (OPT-Meta) and our models, note that setting $\Psi_i(a) = 1$ reduces (OPT-Meta) to the single-firm model (i.e., (Benchmark); see Section 2). Similarly, with a slight abuse of notation, setting $\Psi_i(a) = \Psi_i(a \mid \boldsymbol{a}_1^*, s_1^*)$ (see Equation (13)) reduces (OPT-Meta) to Firm 2's best response (i.e., $\mathsf{OPT}_2(\boldsymbol{a}_1^*, s_1^*)$; see Section 4.1). Note that imposing a non-negativity condition $s \geq 0$ is without loss of optimality, as hiring candidates with negative match values is always suboptimal.

In what follows, we characterize the optimal solution of (OPT-Meta) under Assumption EC.1. Thus, when applying the meta results to derive the structural results of our main models (Appendices EC.3.2 and EC.4.1), much of our proof will focus on verifying that each firm's optimization problem satisfies Assumption EC.1.

ASSUMPTION EC.1. *The functions g_i, h_i, and Ψ_i in (OPT-Meta) satisfy the following:*

(i) Functions g_i and h_i are twice-differentiable and strictly positive density functions

(ii) Function Ψ_i is strictly positive and twice differentiable except for at most one point of jump discontinuity d_i (we set $d_i = \infty$ if such discontinuity point does not exist)

(iii) For each group $i \in \{m, w\}$ and $s \geq 0$, function $F_i(a, s) := \mathbb{E}[(v - s)_+ \mid a, i]\Psi_i(a)$ is increasing in a for $a < d_i$ and has at most one mode for $a > d_i$.

The following propositions establish the optimality conditions and uniqueness of the solution to (OPT-Meta) under Assumption EC.1.

PROPOSITION EC.1 **(Optimality Condition)**. *Under Assumption EC.1, any optimal solution $B^* = (B_m^*, B_w^*)$ and s^* of (OPT-Meta) must satisfy the following:*

(i) The optimal interview set must be greedy with respect to $F_i(a, s^*) = \mathbb{E}[(v - s^*)_+ \mid a, i]\Psi_i(a)$, with a common level θ across both groups:

$$B_i^* = \{a : F_i(a, s^*) \geq \theta\}, \quad \forall i \in \{m, w\}. \tag{EC.4}$$

where θ is uniquely identified to satisfy the interview capacity constraint (EC.2). Furthermore, B_i^* must be a union of at most two disjoint intervals such that $B_i^* = [b_{1,i}^*, d_i] \cup [b_{2,i}^*, b_{3,i}^*]$ for some end points $b_{1,i}^* \leq d_i \leq b_{2,i}^* \leq b_{3,i}^*$.

(ii) The optimal hiring threshold s^* must satisfy:

$$s^* = \min\{s \geq 0 : \sum_{i \in \{m,w\}} 0.5 \int_{B_i^*} \int_s^\infty g_i(v \mid a) h_i(a) \Psi_i(a) \, \mathrm{d}v \, \mathrm{d}a \leq \Delta\}. \tag{EC.5}$$

PROPOSITION EC.2 **(Uniqueness of Optimal Strategy)**. *Under Assumption EC.1, the optimal solution (B^*, s^*) of (OPT-Meta) uniquely exists.*[12]

[12] By uniqueness of the interview set, we mean uniqueness up to a zero-measure set.

The proofs of Propositions EC.1 and EC.2 are presented in Appendices EC.2.1 and EC.2.2, respectively. A few remarks are in order. First, note that the discontinuity point d_i of the function $\Psi_i(a)$ (Assumption EC.1–(ii)) does not depend on s^*. Second, the fact that B_i^* consists of at most two disjoint intervals—one below d_i and one above d_i (Proposition EC.1-(i))—follows directly from the superlevel set representation in (EC.4) and the structural properties of F_i imposed by Assumption EC.1-(iii). Lastly, we emphasize that the endpoints $\{b_{1,i}, b_{2,i}, b_{3,i}\}$ need not be distinct or finite; their values depend on the shape of F_i. For example, if $F_i(a,s)$ is continuously increasing in a for all s (so that $d_i = \infty$, per Assumption EC.1-(ii)), then B_i^* is a single greedy interval. In this case, the representation in (EC.4) simplifies to a form where $b_{1,i}$ is finite while the other endpoints, $b_{2,i}$ and $b_{3,i}$, are infinite. Thus, in applying this meta-result to our main models (Appendicies EC.3.2 and EC.4.1), we leverage the specific shape of F_i (which depends on the firm's optimization problem) to further refine the values of these endpoints—that is, determining whether some collapse or become infinite.

Remark EC.1 (Beyond Normal Distributions). As we highlighted earlier and will show in Appendices EC.3.2 and EC.4.1, the normal distributions used in our model naturally satisfy Assumption EC.1. In particular, we show that for Firm 2, the function $F_i(a,s)$ is unimodal in a for any s, inducing a non-greedy structure in the upper interval of the interview set (Proposition 4).

However, our structural results extend beyond normal distributions. As illustrative examples, we apply our meta-results to two different models beyond our normal distribution framework. First, in Appendix EC.5.1, we recover the structural results of Vohra and Yoder (2023), where the upper interval of Firm 2's interview set is greedy under certain families of conditional match value distributions (such as the exponential distribution). In contrast, in Appendix EC.5.2, we apply the meta-results to the Gumbel distribution (for conditional match value distributions), a workhorse model in the consumer choice framework. We show that the Gumbel distribution results in the unimodal structure of the function $F_i(a,s)$, which consequently leads to the non-greedy structure of Firm 2's interview set, similar to our model with normal distributions.

EC.2.1. Proof of Proposition EC.1

The optimality condition (EC.5) is straightforward because decreasing hiring threshold s can only increase objective value and hiring mass. As such, for any given interview set B, the firm must set as low hiring threshold as possible subject to the hiring capacity constraint (EC.3).

We now prove the first optimality condition (EC.4). Let (B^*, s^*) denote a pair of the optimal interview set $B^* = (B_m^*, B_w^*)$ and hiring threshold s^*. We assume to the contrary. Suppose the optimal interview set B^* is not greedy with respect to $F_i(a, s^*)$ — that is, it does not satisfy Equation (EC.4). We will show that we can construct another solution $(\tilde{B}, \tilde{s})$ that it is feasible — i.e., the resulting interview mass is C, the hiring mass is as most Δ, and $\tilde{s} \geq 0$ — but achieves a strictly higher objective value compared to the original solution (B^*, s^*), which would be a contradiction. In the following, for brevity, we often refer to candidate from group i with score a as candidate (a,i).

We first make the following observation. If $B^* = (B_m^*, B_w^*)$ is not greedy with respect to $F_i(\cdot, s^*)$, it must skip some of the group i candidates with "higher" values of $F_i(\cdot, s^*)$ but interview group j candidates (note that not necessarily $j \neq i$) with "lower" values of $F_j(\cdot, s^*)$. Formally, there exist $\theta \geq 0$ and $i, j \in \{m, w\}$ such that set B_i^* excludes some candidates (a, i) with $F_i(a, s^*) \geq \theta$, but set B_j^* includes some of the candidates (a, j) with $F_j(a, s^*) < \theta$. Because of Assumption EC.1-(iii), the superlevel and sublevel set of $F_i(\cdot, s^*)$ must be a union of disjoint intervals. Hence, the earlier observation implies that, there exist open intervals $X_i := (a_i, b_i)$ (for group i) and $Y_j := (a_j, b_j)$ (for group j) such that[13]

$$B_i^* \cap (a_i, b_i) = \emptyset \tag{EC.6}$$

$$(a_j, b_j) \subseteq B_j^* \tag{EC.7}$$

$$\inf_{a \in (a_i, b_i)} F_i(a, s^*) \geq \theta > \sup_{a \in (a_j, b_j)} F_j(a, s^*) \tag{EC.8}$$

Equation (EC.6) indicates that group i candidates with scores $a \in (a_i, b_i)$ are not interviewed. Line (EC.7) indicates that group j candidates with scores in $a \in (a_j, b_j)$ are interviewed. Finally, the inequality (EC.8) means that all candidates in (a_i, b_i) have higher values of F_i than those of F_j in (a_j, b_j).

We claim that by infinitesimally moving mass from Y_j to X_i, and potentially changing hiring threshold s^* properly, we can construct a feasible pair of interview set and hiring threshold that achieves higher objective than the original optimal solution (B^*, s^*). Without loss of generality, we assume $i \neq j$ (the same argument follows even when $i = j$). Define the following new interview set $\tilde{B}(\epsilon) = (\tilde{B}_m(\epsilon), \tilde{B}_w(\epsilon))$ as follows:

$$\begin{aligned} \tilde{B}_i(\epsilon) &= B_i^* \cup [a_i, a_i + \epsilon] \\ \tilde{B}_j(\epsilon) &= B_j^* \setminus [a_j, a_j(\epsilon)] \end{aligned} \tag{EC.9}$$

for infinitesimally small $\epsilon > 0$. Here, with a slight abuse of notation, we implicitly define $a_j(\epsilon)$ such that $[a_i, a_i + \epsilon]$ and $[a_j, a_j(\epsilon)]$ have the same mass:

$$\int_{a_i}^{a_i + \epsilon} h_i(a)\, \mathrm{d}a = \int_{a_j}^{a_j(\epsilon)} h_j(a)\, \mathrm{d}a, \tag{EC.10}$$

for all $\epsilon \geq 0$. Equivalently, we must have

$$h_i(a_i + \epsilon) = h_j(a_j(\epsilon)) \cdot \frac{\mathrm{d}a_j(\epsilon)}{\mathrm{d}\epsilon} \tag{EC.11}$$

with $a_j(0) = a_j$ by definition. Hence, by construction, the interview set $\tilde{B}(\epsilon)$ has the same interview mass as B^*.

[13] In our proof, we implicitly assumed that it is possible to select an interval from any given set with positive probability measure. Although there exist pathological sets with positive measure that may not contain any intervals, such as the fat Cantor set (see, for example, Aliprantis and Burkinshaw (1998)), these sets can be approximated arbitrarily closely by countable unions of open intervals due to the outer regularity of probability measures (see, for example, Theorem 12.3 in Billingsley (1995)). Therefore, we can extend our argument to accommodate these corner cases.

Denote the objective value and hiring mass of any feasible solution (B,s) as follows:

$$V(B,s)=\sum_{i\in\{m,w\}}0.5\int_{B_i}\int_{s}^{\infty}vg_i(v\mid a)h_i(a)\Psi_i(a)\,\mathrm{d}v\,\mathrm{d}a, \tag{EC.12}$$

$$\lambda(B,s)=\sum_{i\in\{m,w\}}0.5\int_{B_i}\int_{s}^{\infty}g_i(v\mid a)h_i(a)\Psi_i(a)\,\mathrm{d}v\,\mathrm{d}a. \tag{EC.13}$$

Because $\tilde{B}(\epsilon)$ is parameterized by scalar $\epsilon>0$, it is further convenient to define:

$$\tilde{V}(\epsilon,s):=V(\tilde{B}(\epsilon),s), \tag{EC.14}$$

$$\tilde{\lambda}(\epsilon,s):=\lambda(\tilde{B}(\epsilon),s) \tag{EC.15}$$

Note that $\tilde{V}(0,s^*)=V(B^*,s^*)$ and $\tilde{\lambda}(0,s^*)=\lambda(B^*,s^*)$ are the objective value and hiring mass achieved by the original optimal solution (B^*,s^*), respectively.

Given the above notation, we now construct another hiring threshold $\tilde{s}(\epsilon)\geq 0$ such that $(\tilde{B}(\epsilon),\tilde{s}(\epsilon))$ is feasible but achieves a higher objective than (B^*,s^*) for sufficiently small ϵ. To that end, we first observe that, due to Assumption EC.1-(i) and (ii), the second-order partial derivative of $\tilde{V}$ and $\tilde{\lambda}$ exists. In particular,

$$\frac{\partial\tilde{\lambda}(\epsilon,s^*)}{\partial\epsilon}\Big|_{\epsilon=0} \tag{EC.16}$$

is the rate of increases of hiring mass when we change *only* the interview set from B^* to $B(\epsilon)$ with sufficiently small $\epsilon>0$, *while keeping the hiring threshold* s^* *fixed*. With this observation in place, in order to find a new hiring threshold $\tilde{s}(\epsilon)$, we now consider two cases:

Case 1: $s^*=0$ *and* $\dfrac{\partial\tilde{\lambda}(\epsilon,0)}{\partial\epsilon}\Big|_{\epsilon=0}<0$.

In this case, the original hiring threshold is at the lowest value ($s^*=0$), and infinitesimally moving mass from Y_j to X_i can only decrease a total hiring mass ($\dfrac{\partial\tilde{\lambda}(\epsilon,0)}{\partial\epsilon}\Big|_{\epsilon=0}<0$.). Based on this observation, we set $\tilde{s}(\epsilon)=s^*=0$ (i.e., we use the same hiring threshold). Then, for small enough ϵ, we have $\lambda(\tilde{B}(\epsilon),0)=\tilde{\lambda}(\epsilon,0)\leq\tilde{\lambda}(0,0)=\lambda(B^*,0)$ where the inequality holds because $\frac{\tilde{\lambda}(\epsilon,s^*)}{\partial\epsilon}\Big|_{\epsilon=0}<0$. Hence, $(\tilde{B}(\epsilon),0)$ is a feasible pair of interview set and hiring threshold (in the sense that it respects the interview capacity and hiring capacity constraints).

We now investigate whether the objective increases as we switch from $(B^*,0)$ to $(\tilde{B}(\epsilon),0)$. By Taylor's theorem—which can be applied due to Assumption EC.1- (i) and (ii)—we have:[14]

$$\tilde{V}(\epsilon,0)-\tilde{V}(0,0)=\epsilon\cdot\left(\frac{\partial\tilde{V}(\epsilon,0)}{\partial\epsilon}\Big|_{\epsilon=0}\right)+o(\epsilon)=0.5\epsilon h_i(a_i)[F_i(a_i,0)-F_j(a_j,0)]+o(\epsilon). \tag{EC.17}$$

We obtained the last equality in (EC.17) as follows. Observe that, by writing $v=(v-s)+s$, we have:

$$\tilde{V}(\epsilon,s)=s\cdot\tilde{\lambda}(\epsilon,s)+\sum_{i\in\{m,w\}}0.5\int_{\tilde{B}_i(\epsilon)}F_i(a,s)h_i(a).\,\mathrm{d}a \tag{EC.18}$$

[14] For functions $f,g:\mathbb{R}\to\mathbb{R}$, we write $f(x)=o\left(g(x)\right)$ if $\frac{f(x)}{g(x)}\to 0$ as $x\to\infty$.

Hence, taking partial derivative of the above equation with respect to ϵ, we obtain:

$$\frac{\partial \tilde{V}(\epsilon, s^*)}{\partial \epsilon} = \frac{\partial}{\partial \epsilon}\left(s^* \cdot \tilde{\lambda}(\epsilon, s^*) + \sum_{i \in \{m,w\}} 0.5 \int_{\tilde{B}_i(\epsilon)} F_i(a, s^*) h_i(a)\, \mathrm{d}a \right) \tag{EC.19}$$

$$= s^* \cdot \frac{\partial \tilde{\lambda}(\epsilon, s^*)}{\partial \epsilon} + 0.5 \cdot \frac{\partial}{\partial \epsilon}\left(\int_{a_i}^{a_i+\epsilon} F_i(a, s^*) h_i(a)\, \mathrm{d}a - \int_{a_j}^{a_j(\epsilon)} F_j(a, s^*) h_j(a)\, \mathrm{d}a \right) \tag{EC.20}$$

$$= s^* \cdot \frac{\partial \tilde{\lambda}(\epsilon, s^*)}{\partial \epsilon} + 0.5 \cdot h_i(a_i + \epsilon) \cdot [F_i(a_i + \epsilon, s^*) - F_j(a_j(\epsilon), s^*)]. \tag{EC.21}$$

The second line is due to the definition of $\tilde{B}(\epsilon)$ in equation (EC.9). The last line follows from equation (EC.11). Hence, the last equation in (EC.17) follows from evaluating (EC.21) at $\epsilon = 0$ and $s^* = 0$ (recall that $a_j(0) = a_j$ by construction). However, by the construction of a_i and a_j (see inequality (EC.8)) and given $s^* = 0$, we must have $F_i(a_i, 0) > F_j(a_j, 0)$. Thus, for small enough $\epsilon > 0$, we must have $\tilde{V}(\epsilon, 0) - \tilde{V}(0, 0) > 0$. In other words, $(\tilde{B}(\epsilon), 0)$ is feasible and has higher objective than the original optimal solution $(B^*, 0)$, a contradiction.

Case 2: $s^* > 0$ or $\left.\frac{\partial \tilde{\lambda}(\epsilon, s^*)}{\partial \epsilon}\right|_{\epsilon=0} \geq 0$.

In this case, changing the interview set from B^* to $\tilde{B}(\epsilon)$ without adjusting the hiring threshold may increase the total hiring mass. Thus, for a given $\epsilon > 0$, we set the hiring threshold as $\tilde{s}(\epsilon)$, which is defined as the solution to the following equation (with respect to s and parameterized by ϵ):

$$\tilde{\lambda}(\epsilon, s) = \lambda(0, s^*). \tag{EC.22}$$

That is, we set $\tilde{s}(\epsilon)$ such that the hiring mass of $(\tilde{B}(\epsilon), \tilde{s}(\epsilon))$ is equal to that of the original solution (B^*, s^*). Note that $\tilde{s}(0) = s^*$ by construction. By the implicit function theorem, which can be applied due to Assumption EC.1-(i) and (ii), we have:

$$\left.\frac{\mathrm{d}\tilde{s}(\epsilon)}{\mathrm{d}\epsilon}\right|_{\epsilon=0} = -\frac{\left.\frac{\partial \tilde{\lambda}(\epsilon, s^*)}{\partial \epsilon}\right|_{\epsilon=0}}{\left.\frac{\partial \tilde{\lambda}(0, s)}{\partial s}\right|_{s=s^*}}. \tag{EC.23}$$

Note that the denominator is strictly negative (because all g, h, Ψ are strictly positive by Assumption EC.1). Thus, the derivative $\left.\frac{\mathrm{d}\tilde{s}(\epsilon)}{\mathrm{d}\epsilon}\right|_{\epsilon=0}$ is well-defined.

We first show that, for small enough $\epsilon > 0$, $(\tilde{B}(\epsilon), \tilde{s}(\epsilon))$ is a feasible pair of interview set and hiring threshold. Note that, by construction, the pair satisfies the capacity constraints (EC.2) and (EC.3). Thus, it only remains to show that $\tilde{s}(\epsilon) \geq 0$. If $s^* > 0$, the claim is straightforward because $\tilde{s}(0) = s^* > 0$ and $\tilde{s}(\epsilon)$ is differentiable near $\epsilon = 0$ (by the above implicit differentiation). Otherwise, by the hypothesis of case 2, we must have $\left.\frac{\partial \tilde{\lambda}(\epsilon, s^*)}{\partial \epsilon}\right|_{\epsilon=0} \geq 0$, which implies $\left.\frac{\mathrm{d}\tilde{s}(\epsilon)}{\mathrm{d}\epsilon}\right|_{\epsilon=0} \geq 0$ because numerator in Equation (EC.23) is nonnegative and the denominator in Equation (EC.23) is strictly negative. Combining, for sufficiently small $\epsilon > 0$, we have $\tilde{s}(\epsilon) \geq s^* \geq 0$, as desired.

Having established feasibility of $(\tilde{B}(\epsilon), \tilde{s}(\epsilon))$, we now investigate whether the objective increases as we switch from (B^*, s^*) to $(\tilde{B}(\epsilon), \tilde{s}(\epsilon))$. By Taylor's theorem, the change of the objective is given by:

$$\tilde{V}(\epsilon, \tilde{s}(\epsilon)) - \tilde{V}(0, s^*) = \tilde{V}(\epsilon, \tilde{s}(\epsilon)) - \tilde{V}(0, \tilde{s}(0)) \quad (\tilde{s}(0) = s^* \text{ by definition}) \tag{EC.24}$$

$$= \epsilon \cdot \frac{\mathrm{d}\tilde{V}(\epsilon, \tilde{s}(\epsilon))}{\mathrm{d}\epsilon}\Big|_{\epsilon=0} + o(\epsilon). \tag{EC.25}$$

To show that $\tilde{V}(\epsilon, \tilde{s}(\epsilon)) - \tilde{V}(0, s^*) > 0$ for small enough $\epsilon > 0$, it suffices to show that the derivative $\frac{\mathrm{d}\tilde{V}(\epsilon, \tilde{s}(\epsilon))}{\mathrm{d}\epsilon}\Big|_{\epsilon=0}$ is strictly positive. Given $\tilde{s}(0) = s^*$ and using the clain rule, we have:

$$\frac{\mathrm{d}\tilde{V}(\epsilon, \tilde{s}(\epsilon))}{\mathrm{d}\epsilon}\Big|_{\epsilon=0} = \frac{\partial \tilde{V}(\epsilon, s^*)}{\partial \epsilon}\Big|_{\epsilon=0} + \frac{\partial \tilde{V}(0, s)}{\partial s}\Big|_{s=s^*} \cdot \frac{\mathrm{d}\tilde{s}(\epsilon)}{\mathrm{d}\epsilon}\Big|_{\epsilon=0}. \tag{EC.26}$$

We compute each partial derivative as follows. Using Equation (EC.21), we have

$$\frac{\partial \tilde{V}(\epsilon, s^*)}{\partial \epsilon}\Big|_{\epsilon=0} = s^* \cdot \frac{\partial \tilde{\lambda}(\epsilon, s^*)}{\partial \epsilon}\Big|_{\epsilon=0} + 0.5 h_i(a_i)(F_i(a_i, s^*) - F_j(a_j, s^*)). \tag{EC.27}$$

In addition, by taking the partial derivative of line (EC.18) with respect to s, we obtain:

$$\frac{\partial \tilde{V}(0, s)}{\partial s}\Big|_{s=s^*} = \tilde{\lambda}(0, s^*) + s^* \cdot \frac{\partial \tilde{\lambda}(0, s)}{\partial s}\Big|_{s=s^*} + \sum_{i \in \{m,w\}} 0.5 \int_{B_i^*} \left(\frac{\partial F_i(a, s)}{\partial s}\Big|_{s=s^*} \right) h_i(a)\, \mathrm{d}a \tag{EC.28}$$

$$= s^* \cdot \frac{\partial \tilde{\lambda}(0, s)}{\partial s}\Big|_{s=s^*}, \tag{EC.29}$$

where the line (EC.29) follows from the fact that

$$\frac{\partial F_i(a, s)}{\partial s} = -\bar{G}_i(s|a)\Psi_i(a) \tag{EC.30}$$

and therefore

$$\sum_{i \in \{m,w\}} 0.5 \int_{B_i^*} \left(\frac{\partial F_i(a, s)}{\partial s}\Big|_{s=s^*} \right) h_i(a)\, \mathrm{d}a = -\tilde{\lambda}(0, s^*). \tag{EC.31}$$

Substituting (EC.27) and (EC.29) into (EC.26), we obtain

$$\frac{\mathrm{d}\tilde{V}(\epsilon, \tilde{s}(\epsilon))}{\mathrm{d}\epsilon}\Big|_{\epsilon=0} = s^* \cdot \left(\frac{\partial \tilde{\lambda}(\epsilon, s^*)}{\partial \epsilon}\Big|_{\epsilon=0} + \frac{\partial \tilde{\lambda}(0, s)}{\partial s}\Big|_{s=s^*} \cdot \frac{\partial \tilde{s}(\epsilon)}{\mathrm{d}\epsilon}\Big|_{\epsilon=0} \right) + 0.5 h_i(a_i)(F_i(a_i, s^*) - F_j(a_j, s^*)) \tag{EC.32}$$

$$= 0.5 h_i(a_i)(F_i(a_i, s^*) - F_j(a_j, s^*)), \tag{EC.33}$$

where the last line follows from equation (EC.23).

Substituting (EC.33) into (EC.25), we deduce that, for small enough $\epsilon > 0$,

$$\tilde{V}(\epsilon, \tilde{s}(\epsilon)) - \tilde{V}(0, s^*) = 0.5\epsilon h_i(a_i)(F_i(a_i, s^*) - F_j(a_j, s^*)) + o(\epsilon) > 0, \tag{EC.34}$$

where the last inequality is because, by construction of of (a_i, a_j) (see inequality (EC.8)), we must have $F_i(a_i, s^*) > F_j(a_j, s^*)$. Hence, for sufficiently small $\epsilon > 0$, the objective value of $(\tilde{B}(\epsilon), \tilde{s}(\epsilon))$ is strictly higher than that of (B^*, s^*), which is a contradiction.

Thus, we have shown that the optimal interview set B^* and hiring threshold s^* must satisfy

$$B_i^* = \{a : F_i(a, s^*) \geq \theta\}, \quad \forall i \in \{m, w\}. \tag{EC.35}$$

Finally, due to Assumption EC.1-(ii) and (iii), a superlevel set of $F_i(a, s^*)$ (in a) must be a union of at most two intervals: one below and the other above the discontinuity point d_i of $F_i(a, s^*)$. Furthermore, because $F_i(a, s)$ is increasing for $a < d_i$, the lower interval must be of the form $[b_{1,i}, d_i]$. Thus, we conclude that $B_i^* = [b_{1,i}^*, d_i] \cup [b_{2,i}^*, b_{3,i}^*]$, with $b_{1,i}^* \leq d_i \leq b_{2,i}^* \leq b_{3,i}^*$. This completes the proof.

EC.2.2. Proof of Proposition EC.2

We prove that the optimality conditions in in Proposition EC.1 (i.e. equations (EC.4) and (EC.5)) uniquely exists. We proceed in two steps. In Step 1, we express the optimality conditions in Proposition EC.1 as a univariate fixed point equation with respect to a hiring threshold. In Step 2, we show that the fixed point of such equation uniquely exists.

Step 1. Express optimality conditions in Proposition EC.1 as a fixed-point equation. We first provide Definitions EC.1 and EC.2 to simplify the notation of optimality conditions (EC.1) and (EC.5), respectively.

Definition EC.1. *For any given $s \in \mathbb{R}$, define $B(s) = (B_m(s), B_w(s))$ as*

$$B_i(s) = \{a \in \mathbb{R} : F_i(a, s) \geq \theta(s)\}, \tag{EC.36}$$

where level $\theta(s)$ is uniquely identified by the interview capacity constraints. That is, $\theta(s)$ is the unique solution of the following equation:

$$\sum_{i \in \{m,w\}} 0.5 \int_{\{a : F_i(a,s) \geq \theta(s)\}} h_i(a)\, \mathrm{d}a = C. \tag{EC.37}$$

For any $s \geq 0$ and $C \in (0, 1)$, set $B_i(s)$ always uniquely exists. This is because because $h_i(a)$ is a strictly positive and continuous density function and the level $\theta(s)$ is strictly decreases in C. Furthermore, the end points of $B(s)$ are continuous in s due to Assumption EC.1-(i) and (ii).

Definition EC.2. *For a given pair of $(s, \tilde{s})$, we define $\lambda(B(s), \tilde{s})$ as the hiring mass achieved by interview set $B(s)$ given by Definition EC.1 and hiring threshold $\tilde{s}$:*

$$\lambda(B(s), \tilde{s}) = \sum_{i \in \{m,w\}} 0.5 \int_{B_i(s)} \int_{\tilde{s}}^{\infty} g_i(v|a) h_i(a) \Psi_i(a)\, \mathrm{d}v\, \mathrm{d}a. \tag{EC.38}$$

Using Definitions EC.1 and EC.2, we re-write optimality conditions (EC.4) and (EC.5) in Proposition EC.1 as:

$$s^* = \min\{\tilde{s} \geq 0 : \lambda(B(s^*), \tilde{s}) \leq \Delta\}. \tag{EC.39}$$

Note that for any fixed s, function $\lambda(B(s),\tilde{s})$ is strictly decreasing in $\tilde{s}$. Hence, its inverse $\lambda^{-1}(B(s),\cdot)$ is well-defined and must satisfy:

$$\lambda(B(s),\tilde{s}) \le \Delta \Longleftrightarrow \lambda^{-1}(B(s),\Delta) \le \tilde{s}. \tag{EC.40}$$

Combining (EC.39) and (EC.40), we have $s^* = \min\{\tilde{s} \ge 0 : \lambda^{-1}(B(s^*),\Delta) \le \tilde{s}\}$, which implies:

$$\boxed{s^* = \max\{0, \lambda^{-1}(B(s^*),\Delta)\}}. \tag{EC.41}$$

Step 2. Show that the solution of equation (EC.41) uniquely exists. Because $B(s)$ is uniquely pinned down for any given s, it is sufficient to demonstrate that a unique s^* satisfy Equation (EC.41). We establish this uniqueness through the following lemma.

LEMMA EC.1. *Function $\lambda(B(s),s)$ strictly decreases in $s \in \mathbb{R}$.*

Before proving the proof of Lemma EC.1, we first show that Lemma EC.1 implies the existence and uniqueness of a fixed point for equation (EC.41). We first address uniqueness. Assume to the contrary that there exists s^* and $\hat{s}$ satisfying Equation (EC.41) such that $s^* < \hat{s}$. By Lemma EC.1, we have $\lambda(B(s^*),s^*) > \lambda(B(\hat{s}),\hat{s})$. We now consider two cases. First, if $s^* > 0$, then $\hat{s} > 0$ as well. Consequently,

$$s^* = \lambda^{-1}(B(s^*),\Delta) \Longleftrightarrow \lambda(B(s^*),s^*) = \Delta \tag{EC.42}$$

$$\hat{s} = \lambda^{-1}(B(\hat{s}),\Delta) \Longleftrightarrow \lambda(B(\hat{s}),\hat{s}) = \Delta, \tag{EC.43}$$

Thus, we deduce that $\lambda(B(s^*),s^*) = \lambda(B(\hat{s}),\hat{s})$, a contradiction to $\lambda(B(s^*),s^*) > \lambda(B(\hat{s}),\hat{s})$. If $s^* = 0$, using Equation (EC.41), we have $\lambda^{-1}(B(0),\Delta) \le 0$, or equivalently $\lambda(B(0),0) \le \Delta$ from (EC.40). However, we have assumed $\hat{s} > s^* = 0$, and hence (EC.43) holds such that $\lambda(B(0),0) \le \Delta = \lambda(B(\hat{s}),\hat{s})$, which again contradicts Lemma EC.1.

We next address existence. Note that $\lim_{s\to-\infty} \lambda(B(s),s) = C \ge \Delta$ and $\lim_{s\to\infty} \lambda(B(s),s) = 0 < \Delta$. Moreover, the end points of $B_i(s)$ are continuous in s, and therefore $\lambda(B(s),s)$ is continuous in s. Thus, by intermediate value theorem, there exists $s' \in \mathbb{R}$ such that $\lambda(B(s'),s') = \Delta$. If $s' \ge 0$, then $\lambda^{-1}(B(s'),\Delta) = s' \ge 0$ and therefore s' satisfies (EC.41). If $s' < 0$, Lemma EC.1 implies that $\lambda(B(0),0) < \Delta$. Consequently, we have $\lambda^{-1}(B(0),\Delta) < 0$, which implies that $s = 0$ satisfies (EC.41).

Proof of Lemma EC.1. To facilitate subsequent proofs, we begin by expressing the superlevel set $B_i(s)$ as a union of disjoint intervals. Recall from Assumption EC.1-(ii) that the function $\Psi_i(a)$ is discontinuous at most at one point, denoted by d_i (with $d_i = \infty$ if no such point exists). Note that d_i does not depend on s. Furthermore, from Assumption EC.1-(iii), for any $s \ge 0$, function $F_i(a,s)$ increases for $a < d_i$ and has at most one mode at $a > d_i$. Thus, we can express the set $B_i(s)$ as

$$B_i(s) = [b_{1,i}, d_i] \cup [b_{2,i}, b_{3,i}], \tag{EC.44}$$

for some $b_{1,i} \le d_i \le b_{2,i} \le b_{3,i}$. Note that the endpoints $\{b_{1,i}, b_{2,i}, b_{3,i}\}$ are a function of s. Specifically, if endpoint $b_{k,i}$ ($k \in \{1,2,3\}$) is finite, it is determined by solving the equation

$$F_i(b_{k,i}, s) = \theta. \tag{EC.45}$$

In what follows, we focus on this case where all endpoints are finite and distinct. Cases with coinciding or infinite endpoints can be addressed by excluding degenerate intervals or considering only finite endpoints, without affecting the overall proof.[15]

When all endpoints are distinct and finite ($b_{1,i} < d_i < b_{2,i} < b_{3,i} < \infty$), the implicit function theorem, which is applicable due Assumption EC.1-(i) and (ii), ensures that the derivatives $\frac{db_{i,k}}{ds}$ are well-defined for $k \in \{1,2,3\}$. Thus, using the chain rule, the derivative of $\lambda(B(s), s)$ is given by:

$$\begin{aligned}
&\frac{\mathrm{d}\lambda(B(s),s)}{\mathrm{d}s} \\
&= \frac{\mathrm{d}}{\mathrm{d}s} \sum_{i\in\{m,w\}} 0.5 \int_{B_i(s)} \int_s^\infty g_i(v \mid a) h_i(a) \Psi_i(a) \,\mathrm{d}v \,\mathrm{d}a \\
&= \frac{\mathrm{d}}{\mathrm{d}s} \sum_{i\in\{m,w\}} 0.5 \left[\int_{b_{1,i}}^{d_i} \int_s^\infty g_i(v \mid a) h_i(a) \Psi_i(a) \,\mathrm{d}v \,\mathrm{d}a + \int_{b_{2,i}}^{b_{3,i}} \int_s^\infty g_i(v \mid a) h_i(a) \Psi_i(a) \,\mathrm{d}v \,\mathrm{d}a \right] \\
&= \underbrace{-0.5 \sum_{i\in\{m,w\}} \left[\int_{b_{1,i}}^{d_i} g_i(s \mid a) h_i(a) \Psi_i(a) \,\mathrm{d}a + \int_{b_{2,i}}^{b_{3,i}} g_i(s \mid a) h_i(a) \Psi_i(a) \,\mathrm{d}v \,\mathrm{d}a \right]}_{(\mathsf{A})>0} + \\
&\quad \underbrace{0.5 \sum_{i\in\{m,w\}} \left[-\bar{G}_i(s \mid b_{1,i}) h_i(b_{1,i}) \Psi_i(b_{1,i}) \frac{\mathrm{d}b_{1,i}}{\mathrm{d}s} - \bar{G}_i(s \mid b_{2,i}) h_i(b_{2,i}) \Psi_i(b_{2,i}) \frac{\mathrm{d}b_{2,i}}{\mathrm{d}s} + \bar{G}_i(s \mid b_{3,i}) h_i(b_{3,i}) \Psi_i(b_{3,i}) \frac{\mathrm{d}b_{3,i}}{\mathrm{d}s} \right]}_{(\mathsf{B})}
\end{aligned}$$

From the last line, because term $(\mathsf{A}) > 0$, it suffices to prove that term $(\mathsf{B}) \le 0$. To simplify term (B), we first define the following notation. We use $[n]$ to denote set $\{1, 2, ..., n\}$ for positive integer n. Let $\mathcal{E} := \{b_{k,i} : k \in [3], i \in \{m, w\}\}$ denote the set of end points. For each $b \in \mathcal{E}$, we use $i(b)$ to denote the group $i \in \{m, w\}$ such that $b \in B_i(s)$. Then, the partial derivative of function $F_{i(b)}(b, s) = \mathbb{E}[(v - s)_+ \mid b, i = i(b)]$ with respect to s is given by:

$$\frac{\partial F_{i(b)}(b,s)}{\partial s} = -\bar{G}_{i(b)}(s \mid b) \Psi_{i(b)}(b) < 0. \tag{EC.46}$$

We next order and re-label the endpoints $b \in \mathcal{E}$ based on the values of $\left| \frac{\partial F_{i(b)}(b,s)}{\partial s} \right| = \bar{G}_{i(b)}(s \mid b) \Psi_{i(b)}(b)$. Specifically, let m_j (for $j \in [6]$) denote the endpoint $b \in \mathcal{E}$ with the j-th largest value of $\bar{G}_{i(b)}(s \mid b) \Psi_{i(b)}(b)$. In other words, for all $m_j \in \mathcal{E}, j \in [6]$, we have:

$$\bar{G}_{i(m_1)}(s \mid m_1) \Psi_{i(m_1)}(m_1) \ge \bar{G}_{i(m_2)}(s \mid m_2) \Psi_{i(m_2)}(m_2) \ge ... \ge \bar{G}_{i(m_6)}(s \mid m_6) \Psi_{i(m_6)}(m_6), \tag{EC.47}$$

[15] We remind that not all endpoints are required to be finite. For example, if $F_i(a, s)$ is continuous and strictly increasing for all $a \in \mathbb{R}$, which is Firm 1's case (see Appendix EC.3.2), its superlevel set $B_i(s)$ simplifies to $[b_{1,i}, \infty)$, corresponding to setting $d_i = \infty$.

or equivalently,

$$\frac{\partial F_{i(m_1)}(m_1,s)}{\partial s} \le \frac{\partial F_{i(m_2)}(m_2,s)}{\partial s} \le \ldots \le \frac{\partial F_{i(m_6)}(m_6,s)}{\partial s} \le 0. \tag{EC.48}$$

We next define ϵ_j as the change of mass at the end point m_j as we infinitesimally increase s:

$$\epsilon_j := h_{i(m_j)}(m_j) \cdot \left(\frac{\mathrm{d}m_j}{\mathrm{d}s}\right) \cdot (-1)^{\mathbb{1}[m_j \in \mathcal{L}]}, \tag{EC.49}$$

where we use $\mathcal{L}$ to denote a set of *lower* end points, i.e., $\mathcal{L} = \{(b_{1,i}, b_{2,i}) : i \in \{m, w\}\}$. Note that, because $B(s)$ have the identical mass C for all s (Definition EC.1), we must have:

$$\sum_{j=1}^{6} \epsilon_j = 0. \tag{EC.50}$$

With this notation, we can express term (B) as follows:

$$(\mathsf{B}) = \sum_{j=1}^{6} \left| \frac{\partial F_{i(m_j)}(m_j,s)}{\partial s} \right| \cdot \epsilon_j. \tag{EC.51}$$

To prove $(\mathsf{B}) \le 0$, we establish a "cut-off" structure on the sign of ϵ_j through the following claim, which we prove at the end of this section.

Claim EC.1 **(Sign Cutoff of of ϵ_j).** *There exists $l \in [6]$ such that $\epsilon_j \le 0$ if and only if $j \le l$.*

Note that Claim EC.1 directly implies $(\mathsf{B}) \le 0$:

$$(\mathsf{B}) = \sum_{j=1}^{l} \left| \frac{\partial F_{i(m_j)}(m_j,s)}{\partial s} \right| \epsilon_j + \sum_{j=l+1}^{6} \left| \frac{\partial F_{i(m_j)}(m_j,s)}{\partial s} \right| \epsilon_j \tag{EC.52}$$

$$\le \sum_{j=1}^{l} \left| \frac{\partial F_{i(m_l)}(m_l,s)}{\partial s} \right| \epsilon_j + \sum_{j=l+1}^{6} \left| \frac{\partial F_{i(m_l)}(m_l,s)}{\partial s} \right| \epsilon_j \tag{EC.53}$$

$$= 0, \tag{EC.54}$$

where (EC.53) follows from Claim EC.1, along with the definition of m_j's (inequality (EC.47)), and (EC.54) follows from (EC.50).

Proof of Claim EC.1. We recall from equation (EC.45) that, for a given s, the end point m_j is defined through the following equation

$$F_{i(m_j)}(m_j,s) = \theta, \quad \forall j \in [6]. \tag{EC.55}$$

By differentiating with respect to s, we have:

$$\frac{\mathrm{d}m_j}{\mathrm{d}s} = \frac{\dfrac{\partial \theta}{\partial s} - \dfrac{\partial F_{i(m_j)}(m_j,s)}{\partial s}}{\dfrac{\partial F_{i(m_j)}(m_j,s)}{\partial m_j}}. \tag{EC.56}$$

Moreover, because $B_i(s)$ is a superlevel set of $F_i(\cdot, s)$, the lower and upper endpoints of $B_i(s)$ are located in the increasing and decreasing segments of $F_i(\cdot, s)$, respectively. In other words,

$$\frac{\partial F_{i(m_j)}(m_j,s)}{\partial m_j} > 0 \Leftrightarrow m_j \in \mathcal{L}, \tag{EC.57}$$

where we recall that $\mathcal{L} = \{(b_{1,i}, b_{2,i}) : i \in \{m,w\}\}$. Combining (EC.49), (EC.56), and (EC.57), we deduce that

$$\text{sign}(\epsilon_j) = \text{sign}\left(\frac{\frac{\partial \theta}{\partial s} - \frac{\partial F_{i(m_j)}(m_j,s)}{\partial s}}{\frac{\partial F_{i(m_j)}(m_j,s)}{\partial m_j} \cdot (-1)^{\mathbb{1}[m_j \in \mathcal{L}]}}\right) = \text{sign}\left(\frac{\partial F_{i(m_j)}(m_j,s)}{\partial s} - \frac{\partial \theta}{\partial s}\right). \tag{EC.58}$$

Inequality (EC.48) implies that, if $\frac{\partial \theta}{\partial s} \geq \frac{\partial F_{i(m_6)}(m_6,s)}{\partial s}$ (resp. $\frac{\partial \theta}{\partial s} \leq \frac{\partial F_{i(m_1)}(m_1,s)}{\partial s}$), then $\epsilon_j < 0$ (resp., $\epsilon_j < 0$) for all $j \in [6]$, which is a contradiction to the equation (EC.50). Hence, there must exist $l \in [6]$ such that

$$\frac{\partial F_{i(m_1)}(m_1,s)}{\partial s} \leq \ldots \leq \frac{\partial F_{i(m_l)}(m_l,s)}{\partial s} \leq \frac{\partial \theta}{\partial s} \leq \frac{\partial F_{i(m_{l+1})}(m_{l+1},s)}{\partial s} \leq \ldots \leq \frac{\partial F_{i(m_6)}(m_6,s)}{\partial s}. \tag{EC.59}$$

Combining (EC.59) and (EC.58), we conclude that there exists $l \in [6]$ such that $\epsilon_j \leq 0$ if and only if $j \leq l$. □

EC.3. Proofs Related to Section 3

EC.3.1. Preliminary Notions and Auxiliary Results

In the following, we prove some useful properties of the (optimal) hiring threshold endogenously given by an interview threshold (Definition EC.6).

Claim EC.2 **(Properties of Optimal Hiring Threshold)**. *For any given $\boldsymbol{a} = (a_m, a_w)$, define a function $s(\boldsymbol{a})$ as*

$$s(\boldsymbol{a}) := \min\left\{s \geq 0 : \sum_{i \in \{m,w\}} 0.5 \int_{a_i}^{\infty} \int_{s}^{\infty} g_i(v \mid a) h_i(a) \, \mathrm{d}v \, \mathrm{d}a \leq \Delta\right\} \tag{EC.60}$$

Then, $s(\boldsymbol{a})$ satisfy the following properties:

(i) $s(\boldsymbol{a}) = \max(\hat{s}(\boldsymbol{a}), 0)$ for all $\boldsymbol{a} \in \mathbb{R}^2$ where $\hat{s}(\boldsymbol{a})$ is a unique solution of the following equation to s (and parameterized by $\boldsymbol{a}$).[16]

$$\sum_{i \in \{m,w\}} 0.5 \int_{a_i}^{\infty} \int_{s}^{\infty} g_i(v \mid a) h_i(a) \, \mathrm{d}v \, \mathrm{d}a = \Delta. \tag{EC.61}$$

(ii) If $s(\boldsymbol{a}) > 0$, then $\sum_{i \in \{m,w\}} 0.5 \int_{a_i}^{\infty} \int_{s(\boldsymbol{a})}^{\infty} g_i(v \mid a) h_i(a) \, \mathrm{d}v \, \mathrm{d}a = \Delta$.

(iii) $s(\boldsymbol{a})$ is continuous in $\boldsymbol{a}$. Furthermore, its left partial derivative always exists and is non-negative for all values of $\boldsymbol{a}$.

[16] If a solution of the equation does not exist, we set $\hat{s}(\boldsymbol{a}) = -\infty$.

(iv) For any $\boldsymbol{a}$ *such that* $s(\boldsymbol{a})>0$*, the partial derivative of* $s(\boldsymbol{a})$ *exists and is given by*

$$\frac{\partial s(\boldsymbol{a})}{\partial a_i}=-\frac{\int_{s(\boldsymbol{a})}^{\infty} g_i(v\mid a_i)h_i(a_i)\,\mathrm{d}v}{\sum_{i\in\{m,w\}}\int_{a_i}^{\infty} g_i(s(\boldsymbol{a})\mid a)h_i(a)\,\mathrm{d}a}<0,\quad \forall i\in\{m,w\}. \tag{EC.62}$$

Proof of Claim EC.2. Because the hiring mass in equation (EC.60) strictly decreases in s, we have

$$\sum_{i\in\{m,w\}} 0.5\int_{a_i}^{\infty}\int_{s}^{\infty} g_i(v\mid a)h_i(a)\,\mathrm{d}v\,\mathrm{d}a\le\Delta\Leftrightarrow s\ge\hat{s}(\boldsymbol{a}), \tag{EC.63}$$

which directly implies parts (i) and (ii). Furthermore, because the density functions g_i and h_i are continuously differentiable, the implicit function theorem implies that $\hat{s}(\boldsymbol{a})$ is continuously differentiable and its partial derivative is given by

$$\frac{\partial \hat{s}(\boldsymbol{a})}{\partial a_i}=-\frac{\int_{\hat{s}(\boldsymbol{a})}^{\infty} g_i(v\mid a_i)h_i(a_i)\,\mathrm{d}v}{\sum_{i\in\{m,w\}}\int_{a_i}^{\infty} g_i(\hat{s}(\boldsymbol{a})\mid a)h_i(a)\,\mathrm{d}a}<0. \tag{EC.64}$$

Hence, $\hat{s}(\boldsymbol{a})$ is continuous and decreasing in each coordinate. This proves parts (iii) and (iv). □

In the following, we recall standard notion of stochastic orders and related properties (see Shaked and Shanthikumar (2007)). For a univariate random variable X, we use g_X and G_X to denote its p.d.f. and C.D.F., respectively.

Definition EC.3. *Random variables* X *said to be smaller than* Y *in the first stochastic order* $(Y\succeq_{st}X)$ *if* $G_X(t)\ge G_Y(t)$ *for all* $t\in\mathbb{R}$*. Equivalently,* $\mathbb{E}[f(Y)]\ge\mathbb{E}[f(X)]$ *for increasing function* f*.*

Definition EC.4. *Random variables* X *said to be smaller than* Y *in the likelihood ratio order* $(Y\succeq_{lr}X)$ *if* $g_Y(t)/g_X(t)$ *increases in* $t\in\mathbb{R}$*.*

Definition EC.5. *Random variables* X *said to be smaller than* Y *in the increasing convex order* $(Y\succeq_{icx}X)$ *if* $\mathbb{E}[f(X)]\le\mathbb{E}[f(Y)]$ *for all increasing convex function* f*.*

Fact EC.1. *If* $Y\succeq_{lr}X$*, then* $Y\succeq_{st}X$*.*

The following is the standard properties of the normal distributions.

Fact EC.2. *Let* $G(\cdot\mid a)$ *denote the C.D.F. of the normal distribution* $\mathcal{N}(a,\tau^2)$*. Then the family of distributions* $\{G(\cdot\mid a):a\in\mathbb{R}\}$ *increases in* a *in the sense of the likelihood ratio order (Definition EC.4), and therefore in the sense of the first stochastic order (Definition EC.3).*

Fact EC.3. *Let* $X\sim\mathcal{N}(\mu,\sigma^2)$ *with* ϕ *(resp.,* Φ*) denoting p.d.f (resp., C.D.F.) of the standard normal distribution. Then, for any* $a\in\mathbb{R}$*,* $\mathbb{E}[X\mid X>a]=\mu+\sigma\frac{\phi(t)}{1-\Phi(t)}$*, where* $t=\frac{a-\mu}{\sigma}$*. Hence,*

$$\mathbb{E}[(X-a)_+]=(1-\Phi(t))(\mu-a)+\sigma\phi(t). \tag{EC.65}$$

EC.3.2. Proof of Proposition 1

We will apply Propositions EC.1 and EC.2 (Appendix EC.2). To do so, we first verify that the firm's optimization problem satisfies Assumption EC.1 (i)–(iii). By setting $\Psi_i(a) = 1$, (OPT-Meta) in Appendix EC.2 reduces to (Benchmark). Thus, it immediately follows that Assumption EC.1-(i) and (ii) are satisfied. For Assumption EC.1-(iii), consider the excess value function for a given s, defined as $F_i(a, s) := \mathbb{E}[(v - s)_+ | a, i]$. From Fact EC.2, for any s, the function $F_i(a, s)$ is increasing in a. This verifies that Assumption EC.1-(iii) holds. Thus, applying Propositions EC.1 and EC.2, we establish the existence and uniqueness of the firm's optimal solution. Specifically, the firm's optimal interview set $A^* = (A_m^*, A_w^*)$ and hiring level s^* are characterized by optimality conditions (EC.4) and (EC.5). In particular, by (EC.4), the optimal interview set A_i^* for group i must be a superlevel set of $F_i(\cdot, s^*)$ at a common level θ:

$$A_i^* = \{a \in \mathbb{R} : F_i(a, s^*) \geq \theta\}, \quad \forall i \in \{m, w\}, \tag{EC.66}$$

where the level θ is uniquely determined by the interview capacity constraint (2). Since $F_i(\cdot, s^*)$ is increasing, the superlevel set takes the form of a single greedy interval, yielding $A_i^* = [a_i^*, \infty)$ for some interview threshold a_i^*. This proves part (a). Part (b) follows from optimality condition (EC.5).

EC.3.3. Proof of Corollary 1

First, building on Proposition 1, we first characterize the firm's optimal interview threshold $\boldsymbol{a}^* = (a_m, a_w^*)$ as a unique solution of Equations (EC.67) and (EC.68).

Lemma EC.2. *The optimal interview threshold (a_m^*, a_w^*) is a unique solution satisfying the following:*

$$\mathbb{E}[(v - s(\boldsymbol{a}^*))_+ | a = a_m^*, m] = \mathbb{E}[(v - s(\boldsymbol{a}^*)_+) | a = a_w^*, w] \tag{EC.67}$$

$$\sum_{i \in \{m,w\}} 0.5 \int_{a_i^*} h_i(a)\, \mathrm{d}a = C \tag{EC.68}$$

where function $s(\boldsymbol{a})$ is defined in (EC.60) (see Claim EC.2 in Appendix EC.3.1).

Proof of Lemma EC.2. The proof directly follows from Proposition 1 and continuity of the normal density functions h_i. □

Fact EC.4 (**Theorem 4 of Müller (2001)**). *Let $X \sim \mathcal{N}(\mu_1, \sigma_1^2)$ and $Y \sim \mathcal{N}(\mu_2, \sigma_2^2)$. Then $Y \succeq_{icx} X$ (see Definition EC.5) if any only if $\mu_1 \leq \mu_2$ and $\sigma_1^2 \leq \sigma_2^2$.*

For any $s \in \mathbb{R}$, note that $(v - s)_+$ is an increasing convex function in v. If $a_m \leq a_w$, because $\tau_m < \tau_w$ (Assumption 1-(b)), Fact EC.4 implies that $\mathbb{E}[(v - s)_+ | a = a_m, m] < \mathbb{E}[(v - s)_+ | a = a_w, w]$ for any $s \in \mathbb{R}$.[17] Hence, any $\boldsymbol{a}^*$ that satisfy equation (EC.67) must satisfy $a_m^* > a_w^*$. This completes the proof.

[17] One can show that the inequality must be strict as long as $\tau_m \neq \tau_w$.

EC.3.4. Proof of Corollary 2

By mirroring the proof of Proposition EC.1, we can establish that the optimal interview set under (Intervention) must be greedy in scores within each group, implying $A_i^\rho = [a_i^\rho, \infty)$ for some interview threshold $\boldsymbol{a}^\rho = (a_m^\rho, a_w^\rho)$. From Proposition 1, the optimal interview threshold $\boldsymbol{a}^*$ under (Benchmark) uniquely exists. If $\boldsymbol{a}^*$ already satisfies the ρ-Rooney Rule, then it remains the unique optimal solution under (Intervention), implying $\boldsymbol{a}^\rho = \boldsymbol{a}^*$. Otherwise, due to the uniqueness of the unconstrained optimal solution under (Benchmark), the ρ-Rooney Rule must bind, determining the interview threshold through the equations: $0.5(1 - H_w(a_w^\rho)) = \rho C$ and $0.5(1 - H_m(a_m^\rho)) = (1-\rho)C$. Given $\boldsymbol{a}^\rho$, the firm then maximizes its objective by hiring as many interviewees as possible with the highest non-negative match values. Hence, the optimal hiring threshold s^ρ satisfies Equation (9). This completes the proof.

EC.3.5. Proof of Proposition 2

Let $\boldsymbol{a}' = (a'_m, a'_w)$ equalizes the interview mass for both groups, or equivalently, $\mathbb{P}[a \geq a'_m \mid m] = \mathbb{P}[a \geq a'_w \mid w]$. Under this condition, we claim the following.

Claim EC.3. *If the finite interview threshold $\boldsymbol{a}'$ satisfies $\mathbb{P}[a \geq a'_m \mid m] = \mathbb{P}[a \geq a'_w \mid w]$, then $\mathbb{P}[v \geq s \mid a \geq a'_m, m] > \mathbb{P}[v \geq s \mid a \geq a'_w, w]$ for any finite s.*

Before proving this claim, we note that the claim directly implies Proposition 2. Specifically, let $g_i(a, v)$ be the joint p.d.f of (a, v) of group i. Claim EC.3 implies that, for any finite s, we have

$$\mathbb{P}[v \geq s \mid a \geq a'_m, m] > \mathbb{P}[v \geq s \mid a \geq a'_w, w] \tag{EC.69}$$

$$\Leftrightarrow \frac{\int_s^\infty \int_{a'_m}^\infty g_m(v \mid a) h_m(a)\, \mathrm{d}a\, \mathrm{d}v}{\mathbb{P}[a \geq a'_m \mid m]} > \frac{\int_s^\infty \int_{a'_w}^\infty g_w(v \mid a) h_w(a)\, \mathrm{d}a\, \mathrm{d}v}{\mathbb{P}[a \geq a'_w \mid w]} \tag{EC.70}$$

$$\overset{(*)}{\Leftrightarrow} \int_s^\infty \int_{a'_m}^\infty g_m(v \mid a) h_m(a)\, \mathrm{d}a\, \mathrm{d}v > \int_s^\infty \int_{a'_w}^\infty g_w(v \mid a) h_w(a)\, \mathrm{d}a\, \mathrm{d}v \tag{EC.71}$$

where $(*)$ holds because the interview mass is identical between the two groups by our assumption.

It only remains to prove Claim EC.3. For this, we will use the following fact:

Fact EC.5 **(Single Crossing of p.d.f. Implies First-order Stochastic Dominance)**. *Let X and Y be random variables with continuous p.d.f g_X and g_Y, respectively, that satisfy $g_X(v) > g_Y(y)$ (resp., $g_X(v) < g_Y(y)$) if $v > \underline{v}$ (resp., $v < \underline{v}$). Then $\mathbb{P}[X \geq s] > \mathbb{P}[Y \geq s]$ for all finite s.*

Proof of Fact EC.5. If $s \leq \underline{v}$, then $\mathbb{P}[X < s] = \int_{-\infty}^s g_X(v)dv < \int_{-\infty}^s g_Y(v)dv = \mathbb{P}[Y < s]$ (unless $s = \infty$). If $s \geq \underline{v}$, then $\mathbb{P}[X \geq s] = \int_s^\infty g_X(v)dv > \int_s^\infty g_Y(v)dv = \mathbb{P}[x \geq s]$ (unless $s = -\infty$). □

To apply Fact EC.5, we will compare p.d.f of v of group i conditional on being interviewed. Specifically, recall that the marginal distribution of v is identical across group $i \in \{m, w\}$ by Assumption 1-(a). Let $\kappa^2 := \sigma_i^2 + \tau_i^2$ denote the common unconditional variance of v of both groups. We further use $g(v)$ to denote the common marginal p.d.f. of v for each group. Then, the conditional p.d.f of v, given that a candidate of group i is interviewed (with $\boldsymbol{a}'$ ensuring equal interview rates), is given by:

$$g(v \mid a \geq a', i) = \frac{g(v) \cdot \mathbb{P}[a \geq a'_i \mid v, i]}{\mathbb{P}[a \geq a'_i \mid i]}. \tag{EC.72}$$

On the other hand, we have $v \mid a, i \sim \mathcal{N}(a, \tau_i^2)$ and $a \mid i \sim \mathcal{N}(0, \sigma_i^2)$. Thus, the posterior distribution of score a of group i conditional on match value v is given by $a \mid \{v, i\} \sim \mathcal{N}\left(\frac{\sigma_i^2 v}{\sigma_i^2+\tau_i^2}, \left(\frac{1}{\sigma_i^2}+\frac{1}{\tau_i^2}\right)^{-1}\right)$.

Given the above two observations, we now show the following:

$$g(v \mid a \geq a'_m, m) > g(v \mid a \geq a'_w, w) \tag{EC.73}$$

$$\Leftrightarrow \mathbb{P}[a \geq a'_m \mid v, m] > \mathbb{P}[a \geq a'_w \mid v, w] \tag{EC.74}$$

$$\Leftrightarrow \Phi^c\left(\left(a'_m - \frac{\sigma_m^2 v}{\sigma_m^2+\tau_m^2}\right)\cdot\sqrt{\frac{1}{\sigma_m^2}+\frac{1}{\tau_m^2}}\right) > \Phi^c\left(\left(a'_w - \frac{\sigma_w^2 v}{\sigma_w^2+\tau_w^2}\right)\cdot\sqrt{\frac{1}{\sigma_w^2}+\frac{1}{\tau_w^2}}\right) \tag{EC.75}$$

$$\Leftrightarrow \left(a'_m - \frac{\sigma_m^2 v}{\kappa^2}\right)\cdot\sqrt{\frac{\kappa^2}{\sigma_m^2\tau_m^2}} < \left(a'_w - \frac{\sigma_w^2 v}{\kappa^2}\right)\cdot\sqrt{\frac{\kappa^2}{\sigma_w^2\tau_w^2}} \tag{EC.76}$$

$$\Leftrightarrow a'_m\sqrt{\frac{\kappa^2}{\sigma_m^2\tau_m^2}} - a'_w\sqrt{\frac{\kappa^2}{\sigma_w^2\tau_w^2}} < \left(\frac{\sigma_m}{\tau_m}-\frac{\sigma_w}{\tau_w}\right)\frac{v}{\kappa}. \tag{EC.77}$$

In (EC.74), we used the assumptions that (i) $\mathbb{P}[a \geq a'_m \mid m] = \mathbb{P}[a \geq a'_w \mid w]$ and (ii) the two groups have common marginal p.d.f. of v as $g(v)$. In (EC.75), we used the posterior distribution of $a \mid \{v, i\}$. Line (EC.76) follows from $\sigma_i^2 + \tau_i^2$ is a common value κ^2 shared by the two groups from Assumption 1-(a). Finally, Assumption 1-(b) implies that $\sigma_m/\tau_m > \sigma_w/\tau_w$. Note that the left-hand side of (EC.77) is finite by the assumption of $\boldsymbol{a}'$ being finite. Hence, by Fact EC.5, we conclude that $\mathbb{P}[v \geq s \mid a \geq a'_m, i = m] > \mathbb{P}[v \geq s \mid a \geq a'_w, i = w]$ for all finite s. This completes the proof.

EC.3.6. Detailed Proof Outline of Theorem 1

In this section, we provide a detailed outline of the proof for Theorem 1. We proceed in three steps.[18]

Step 1: Iso-interview Curve is above Iso-hiring Curve. We first introduce some notation. Consider a scenario where the firm follows a greedy interview strategy with a given interview threshold $\boldsymbol{a} = (\boldsymbol{a}_m, \boldsymbol{a}_w)$. Given this interview threshold, the firm chooses a hiring threshold to maximize its total match value subject to the hiring capacity Δ. Note that, the optimal hiring policy in this scenario is to hire as many interviewees as possible with the highest non-negative match values to fill the hiring capacity. Building on this observation, we define the interview mass and hiring functions as follows.

Definition EC.6 **(Interview & Hiring Mass Given Interview Thresholds)**. *For any given interview threshold $\boldsymbol{a} = (a_m, a_w)$, we define the following functions:*

$$\eta_i(\boldsymbol{a}) := 0.5\int_{a_i} h_i(a)\,\mathrm{d}a \tag{EC.78}$$

$$s(\boldsymbol{a}) := \min\left\{s \geq 0 : \sum_{i\in\{m,w\}} 0.5\int_{a_i}\int_s^\infty g_i(v \mid a)h_i(a)\,\mathrm{d}v\,\mathrm{d}a \leq \Delta\right\} \tag{EC.79}$$

$$\lambda_i(\boldsymbol{a}) := 0.5\int_{a_i}\int_{s(\boldsymbol{a})}^\infty g_i(v \mid a)h_i(a)\,\mathrm{d}v\,\mathrm{d}a \tag{EC.80}$$

[18] Throughout the proof, we implicit assume $C < 1$ because otherwise the optimal interview threshold ill-defined (i.e., the threshold is $-\infty$), and by Assumption 1, both groups are trivially interviewed and hired at equal rates.

Here, $\eta_i(\boldsymbol{a})$ represents the interview mass of group i given the interview threshold $\boldsymbol{a}$. This threshold $\boldsymbol{a}$ subsequently determines the optimal (endogenous) hiring threshold $s(\boldsymbol{a})$ via Equation (EC.79), which in turn defines the hiring mass of group i, denoted by $\lambda_i(\boldsymbol{a})$. Note that, for a given interview capacity C, the optimal interview and hiring mass of group i under (Benchmark) is given by $\eta_i(\boldsymbol{a}^*)$ and $\lambda_i(\boldsymbol{a}^*)$, respectively.[19]

En route to establishing our desired results, we first establish a monotonicity property of the hiring mass function $\lambda_i(\boldsymbol{a})$ in the following lemma.

LEMMA EC.3 **(Monotonicity of the Hiring Mass Function)**. *For each $i \in \{m,w\}$, fixing a_j for $j \neq i$, the function $\lambda_i(\boldsymbol{a})$ strictly decreases in a_i.*

We prove Lemma EC.3 in Appendix EC.3.7. Note that Lemma EC.3 implies that the epigraph of the iso-hiring curve in Figure 2-(b) is the set of interview thresholds where group w is hired less than group m. Given Definition EC.6 and Lemma EC.3, we formally show that the iso-interview curve is above the iso-hiring curve in the following.

PROPOSITION EC.3 **(Iso-interview Curve is Above Iso-hiring Curve)**. *Define sets ζ and γ as:*

$$\begin{aligned} \zeta(a_m) &:= \{a_w \in \mathbb{R} : \lambda_m(\boldsymbol{a}) - \lambda_w(\boldsymbol{a}) = 0\} \\ \gamma(a_m) &:= \{a_w \in \mathbb{R} : \eta_m(\boldsymbol{a}) - \eta_w(\boldsymbol{a}) = 0\} \end{aligned} \tag{EC.81}$$

Then, $\zeta(\cdot)$ and $\gamma(\cdot)$ are functions on $\mathbb{R}$, i.e. $|\zeta(a_m)| = 1$ and $|\gamma(a_m)| = 1$ for all $a_m \in \mathbb{R}$. Furthermore, $\mathsf{epi}(\gamma) \subset \mathsf{epi}(\zeta)$ where $\mathsf{epi}(f)$ is the epigraph of function f.

We prove Proposition EC.3 in Appendix EC.3.8. The proof leverages Proposition 2 which states that, at any interview thresholds on the iso-interview curve, group w is strictly hired less. As such, due to Lemma EC.3, fixing group m's interview thresholds, the only way to close the hiring gap is to decrease group w's interview threshold. Formally, we show that $\zeta(a_m) < \gamma(a_m)$ and thus $\mathsf{epi}(\gamma) \subset \mathsf{epi}(\zeta)$.

Step 2: The Optimal Interview Thresholds Curve Crosses the Iso-interview Curve Only Once. The following proposition formalizes that the optimal interview thresholds curve intersects the iso-interview curve only once. Specifically, the firm finds it optimal to over-represent group w in the interview set if and only if C exceeds a certain threshold.

PROPOSITION EC.4 **(Single Crossing of Optimal Interview Thresholds and Iso-interview Curve)**. *For small enough Δ, there exists a unique interview capacity threshold $\underline{C} \in [\Delta, 1)$ such that:*

$$\eta_m(\boldsymbol{a}^*) = \eta_w(\boldsymbol{a}^*) \quad \text{if and only if} \quad C = \underline{C}.$$

Furthermore, for any limited interview capacity $C \in [\Delta, 1)$, it holds that $\eta_m(\boldsymbol{a}^*) > \eta_w(\boldsymbol{a}^*)$ if and only if $C < \underline{C}$.

[19] We remind that the optimal interview threshold $\boldsymbol{a}^*$ under (Benchmark) is a function of C, but we omit its dependence on C for brevity.

We prove Proposition EC.4 in Appendix EC.3.9. The proof establishes several properties of $\boldsymbol{a}^*$ with respect to interview capacity C: (i) As the firm increases C, it lowers the thresholds a_m^* and a_w^*, expanding interview sets for both groups (Lemma EC.6 - (b)). (ii) However, a_w^* decreases *faster* than a_m^* (Lemma EC.6 - (c)), which formalizes our earlier intuition that a larger interview capacity allows the firm to explore more candidates from group w.

Step 3: Putting everything together. Finally, we combine the previous two steps to prove Theorem 1. Define:

$$\begin{aligned}\underline{C} &:= \{C \in [\Delta, 1) : \eta_m(\boldsymbol{a}^*) = \eta_w(\boldsymbol{a}^*)\},\\ \overline{C} &:= \min\{C \in [\Delta, 1) : \lambda_m(\boldsymbol{a}^*) = \lambda_w(\boldsymbol{a}^*)\}.\end{aligned} \tag{EC.82}$$

That is, $\underline{C}$ is the value of (limited) interview capacity C for which the optimal interview threshold intersects with the iso-interview curve. Note that, by Proposition EC.4, $\underline{C}$ uniquely exists. On the other hand, $\overline{C}$ is the smallest interview capacity such that the optimal interview threshold intersects with the iso-hiring curve.

Now, by Proposition EC.3 and the continuity of $\boldsymbol{a}^*$ in C (see Lemma EC.6-(a)), we must have $\underline{C} < \overline{C}$. With this observation in place to prove Theorem 1, we consider two cases. First, if $C < \underline{C}$, group w is under-represented in both the optimal interview set and the final hiring outcome, as $\boldsymbol{a}^*$ belongs to the epigraph of both the iso-interview and iso-hiring curves (green-colored area in Figure 2). In this case, by Corollary 2, the 0.5-Rooney rule forces the firm to interview more (less, resp.) of group w (group m, resp.). From Lemma EC.3, this adjustment strictly improves group w's hiring mass. However, from Proposition 2, it cannot achieve the equal hiring, thus proving part (a) of Theorem 1. On the other hand, if $C \in [\underline{C}, \overline{C}]$, the optimal interview thresholds belong to the epigraph of the iso-hiring curve but lie below the iso-interview curve. As such, the group w is over-represented in the interview set but under-represented in the hiring outcome (blue-colored area in Figure 2). In this case, again due to Corollary 2, the 0.5-Rooney rule does not change the firm's interview and hiring decisions, leading to part (b) of Theorem 1. This completes the proof.

EC.3.7. Proof of Lemma EC.3

We prove a stronger result in the following (that directly implies Lemma EC.3).

LEMMA EC.4 **(Monotonicity of Hiring Mass in Interview Threshold)**. *Let* $\boldsymbol{a} = (a_m, a_w)$ *and* $\boldsymbol{a}' = (a'_m, a'_w)$ *satisfy* $a'_i \geq a_i$ *and* $a'_j \leq a_j$ *for two distint groups* $i \neq j$*, and at least one of the two inequalities is strict. Then,* $\lambda_i(\boldsymbol{a}') < \lambda_i(\boldsymbol{a})$ *and* $\lambda_j(\boldsymbol{a}') > \lambda_j(\boldsymbol{a})$ *(See Definition EC.6).*

Proof of Lemma EC.4. Without loss of generlaity, we consider the case where $i = m$ and $j = w$. Let $\lambda(\boldsymbol{a}) := \lambda_m(\boldsymbol{a}) + \lambda_w(\boldsymbol{a})$ for any $\boldsymbol{a}$. We further define

$$\begin{aligned}\mathrm{d}s &:= s(\boldsymbol{a}') - s(\boldsymbol{a}),\\ \mathrm{d}\lambda &:= \lambda(\boldsymbol{a}') - \lambda(\boldsymbol{a}).\end{aligned} \tag{EC.83}$$

We consider four cases described in the following table.

	$\mathrm{d}\lambda > 0$	$\mathrm{d}\lambda = 0$	$\mathrm{d}\lambda < 0$
$\mathrm{d}s > 0$	case 1	case 1	case 4
$\mathrm{d}s = 0$	case 2	case 2	case 2
$\mathrm{d}s < 0$	case 4	case 3	case 3

(i) **case 1:** Because the density functions g_i and h_i are strictly positive, we observe that:

$$\lambda_m(\boldsymbol{a}') = 0.5 \int_{a'_m} \int_{s(\boldsymbol{a}')} g_m(v|a) h_m(a)\, \mathrm{d}v\, \mathrm{d}a \tag{EC.84}$$

$$< 0.5 \int_{a_m} \int_{s(\boldsymbol{a})} g_m(v|a) h_m(a)\, \mathrm{d}v\, \mathrm{d}a \tag{EC.85}$$

$$= \lambda_{1,m}(\boldsymbol{a}) \tag{EC.86}$$

where the inequality follows from $\mathrm{d}s > 0$ and $a'_m \geq a_m$, as assumed in Case 1. Hence, $\lambda_m(\boldsymbol{a}') < \lambda_m(\boldsymbol{a})$, which implies that $\lambda_w(\boldsymbol{a}') > \lambda_w(\boldsymbol{a})$ because $\mathrm{d}\lambda \geq 0$.

(ii) **cases 2 and 3**: The proof follows similarly to case 1 and is therefore omitted for brevity.

(iii) **case 4:** We claim that this case cannot occur. Suppose for the sake of contradiction that $\mathrm{d}s > 0$ and $\mathrm{d}\lambda < 0$. Then, either of the following two sub-cases happens. In the first sub-case, $s(\boldsymbol{a}) = 0$ and $s(\boldsymbol{a}') > 0$. However, from Claim EC.2-(ii) (Appendix EC.3.1), this implies that $\lambda(\boldsymbol{a}) \leq \Delta = \lambda(\boldsymbol{a}')$, implying that $\mathrm{d}\lambda \geq 0$. In the second sub-case, both of $s(\boldsymbol{a})$ and $s(\boldsymbol{a}')$ are strictly positive, but again due to Claim EC.2-(ii), this implies that $\lambda(\boldsymbol{a}) = \lambda(\boldsymbol{a}') = \Delta$. In either subcase, we reach to a contradiction to $\mathrm{d}\lambda < 0$. By following a similar argument, one can also show a contradiction for $\mathrm{d}s < 0$ and $\mathrm{d}\lambda > 0$. This completes the proof. □

EC.3.8. Proof of Proposition EC.3

We first characterize $\gamma(\cdot)$. Because score a of group i is normally distributed with mean zero and variance σ_i^2, we deduce that $\eta_{1,m}(\boldsymbol{a}) = \eta_{1,w}(\boldsymbol{a})$ if and only if $\dfrac{a_m}{\sigma_m} = \dfrac{a_w}{\sigma_w}$. In other words, we have $\gamma(a_m) = \dfrac{\sigma_w}{\sigma_m} a_m$. We now turn our attention to $\zeta(\cdot)$. First, we show that it is a function. Note that

$$\lim_{a_w \to \infty} (\lambda_m(\boldsymbol{a}) - \lambda_w(\boldsymbol{a})) > 0.$$

because the hiring mass of group w converges to zero as $a_w \to \infty$. Furthermore, letting $\overline{s} := \lim_{a_w \to -\infty} s(\boldsymbol{a})$ and due to the continuity of $s(\boldsymbol{a})$ (see Claim EC.2 in Appendix EC.3.1), we have:

$$\lim_{a_w \to -\infty} \lambda_w(\boldsymbol{a}) = 0.5 \int_{\overline{s}}^{\infty} \int_{-\infty}^{\infty} g_w(v|a) h_w(a)\, \mathrm{d}a\, \mathrm{d}v = 0.5\, \mathbb{P}_{v \sim \mathcal{N}(0,\kappa^2)}[v \geq \overline{s}]. \tag{EC.87}$$

where $\kappa^2 := \sigma_w^2 + \tau_w^2 = \sigma_m^2 + \tau_m^2$ (Assumption 1-(a)). On the other hand,

$$\lim_{a_w \to -\infty} \lambda_m(\boldsymbol{a}) = 0.5 \int_{a_m}^{\infty} \int_{\overline{s}}^{\infty} g_m(v|a)\, \mathrm{d}v\, \mathrm{d}a < 0.5 \int_{-\infty}^{\infty} \int_{\overline{s}}^{\infty} g_m(v|a)\, \mathrm{d}v\, \mathrm{d}a = 0.5\, \mathbb{P}_{v \sim \mathcal{N}(0,\kappa^2)}[v \geq \overline{s}] \tag{EC.88}$$

where we again used Assumption 1-(a) in the last equality. Hence, we must have

$$\lim_{a_w \to -\infty} (\lambda_m(\boldsymbol{a}) - \lambda_w(\boldsymbol{a})) < 0. \tag{EC.89}$$

Therefore, by the intermediate value theorem, $|\zeta(a_m)| \geq 1$. Furthermore, fixing a_m, Lemma EC.4 implies that $\lambda_m(\boldsymbol{a}) - \lambda_w(\boldsymbol{a})$ is strictly increasing in a_w. Thus, we deduce that $|\zeta(a_m)| = 1$ and therefore $\zeta(\cdot)$ indeed defines a function.

Finally, we show that $\mathsf{epi}(\gamma) \subset \mathsf{epi}(\zeta)$. Because γ and ζ are functions on $\mathbb{R}$, it suffices to show that $\gamma(a_m) > \zeta(a_m)$ for all $a_m \in \mathbb{R}$. From Proposition 2,

$$\lambda_m(a_m, \gamma(a_m)) - \lambda_w(a_m, \gamma(a_m)) > 0. \tag{EC.90}$$

Moreoever, from Lemma EC.4, fixing a_m, the hiring gap $\lambda_m(\boldsymbol{a}) - \lambda_w(\boldsymbol{a})$ increases in a_w. Hence, by definition of $\zeta(a_m)$, we must have $\gamma(a_m) > \zeta(a_m)$. This completes the proof.

EC.3.9. Proof of Proposition EC.4

To prove Proposition EC.4, we study how the optimal interview threshold changes in the interview capacity C. Toward that end, we first establish several auxiliary results. We first recall from Lemma EC.2 (Appendix EC.3.3) that the optimal interview threshold can be obtained by solving a system of equations. For ease of reference, we restate the two equations defining the optimal interview threshold from Lemma EC.2:

$$\mathbb{E}[(v - s(\boldsymbol{a}^*))_+ | a = a_m^*, m] = \mathbb{E}[(v - s(\boldsymbol{a}^*))_+ | a = a_w^*, w] \tag{EC.91}$$

$$\sum_{i \in \{m,w\}} 0.5 \int_{a_i^*} h_i(a)\, \mathrm{d}a = C \tag{EC.92}$$

where function $s(\boldsymbol{a})$ is given by

$$s(\boldsymbol{a}) := \min \left\{ s \geq 0 : \sum_{i \in \{m,w\}} 0.5 \int_{a_i} \int_s^\infty g_i(v|a) h_i(a)\, \mathrm{d}v\, \mathrm{d}a \leq \Delta \right\} \tag{EC.93}$$

We establish some useful properties of the solution that satisfies Equation (EC.91).

Lemma EC.5. *For a given $\boldsymbol{a} = (a_m, a_w)$, define function $f(\boldsymbol{a})$ as follows:*

$$f(\boldsymbol{a}) := \mathbb{E}[(v - s(\boldsymbol{a}))_+ | a = a_m, m] - \mathbb{E}[(v - s(\boldsymbol{a}))_+ | a = a_w, w]. \tag{EC.94}$$

Then, for any $\boldsymbol{a}$ such that $f(\boldsymbol{a}) = 0$, we have:

(a) $\bar{G}_m(s(\boldsymbol{a})|a_m) > \bar{G}_w(s(\boldsymbol{a})|a_w)$

(b) Let $\omega(a_m) := \{a_w : f(\boldsymbol{a}) = 0\}$. Then $\omega(a_m)$ is a strictly increasing function.

Proof of Lemma EC.5-(a). Throughout the proof, we remind that we use ϕ and Φ to denote p.d.f. and C.D.F. of the standard normal distribution, respectively. The following fact will be further useful.

Fact EC.6. *For any random variable X with continuous C.D.F. (p.d.f., resp.) $G(\cdot)$ ($g(\cdot)$, resp.) and $s \in \mathbb{R}$,*

$$\mathbb{E}[(X - s)_+] = \int_s^\infty (1 - G(x)) dx.$$

Proof of Fact EC.6.

$$\begin{aligned}\mathbb{E}[(X-s)_+] &= \int_s^\infty (x-s)g(x)dx = \int_s^\infty xg(x)dx - s(1-G(s)) = \int_s^\infty \int_0^x g(x)dydx - s(1-G(s)) \\ &\overset{(*)}{=} \int_0^s \int_s^\infty g(x)dxdy + \int_s^\infty \int_y^\infty g(x)dxdy - s(1-G(s)) \\ &= \int_0^s (1-G(s))dy + \int_s^\infty (1-G(y))dy - s(1-G(s)) \\ &= \int_s^\infty (1-G(y))dy.\end{aligned}$$

In equality $(*)$, we changed the order of integrals. The rest of the steps are algebraic. □

From Fact EC.6, we deduce that $f(\boldsymbol{a}) = 0$ if and only if:

$$\int_{s(\boldsymbol{a})}^\infty \Phi^c\left(\frac{x-a_m}{\tau_m}\right)dx = \int_{s(\boldsymbol{a})}^\infty \Phi^c\left(\frac{x-a_w}{\tau_w}\right)dx. \tag{EC.95}$$

We now observe that, because $\tau_m < \tau_w$ (by Assumption 1),

$$\Phi^c\left(\frac{x-a_m}{\tau_m}\right) > \Phi^c\left(\frac{x-a_w}{\tau_w}\right) \Leftrightarrow x < \overline{x} := \frac{a_m/\tau_m - a_w/\tau_w}{1/\tau_m - 1/\tau_w}. \tag{EC.96}$$

Hence, any $s(\boldsymbol{a})$ that satisfies (EC.95) must satisfy $s(\boldsymbol{a}) < \overline{x}$, or equivalently,

$$\Phi^c\left(\frac{s(\boldsymbol{a})-a_m}{\tau_m}\right) > \Phi^c\left(\frac{s(\boldsymbol{a})-a_w}{\tau_w}\right), \tag{EC.97}$$

which leads to the desired result. This completes the proof. □

Proof of Lemma EC.5-(b). In what follows, we use $\frac{\partial^- f(\mathbf{a})}{\partial a_i}$ to denote the left partial derivative of function f with respect to a_i. Note that, from claim EC.2, the left partial derivative of $f(\boldsymbol{a})$ exists because that of $s(\boldsymbol{a})$ exists.

We proceed in two steps. First, in Claim EC.4, we show that when a_w is fixed, there is a unique a_m such that $f(\boldsymbol{a}) = 0$. This does not, by itself, guarantee uniqueness in the other direction (fixing a_m and solving for a_w). However, by analyzing how the sign of f changes near its root, we show that uniqueness in the other direction also holds. Such "bi-directional" uniqueness will readily imply that the curve $\{\boldsymbol{a} \in \mathbb{R}^2 : f(\boldsymbol{a}) = 0\}$ defines an injective function, implying that $\omega(a_m)$ is strictly monotone.

Claim EC.4. *Define $\beta(a_w) := \{a_m : f(\boldsymbol{a}) = 0\}$. Then, for any $\beta \in \beta(a_w)$,*

$$\frac{\partial^- f(\boldsymbol{a})}{\partial a_m}\Big|_{(\beta, a_w)} > 0. \tag{EC.98}$$

Furthermore, $|\beta(a_w)| = 1$ for any $a_w \in \mathbb{R}$.

Proof of Claim EC.4. From Fact EC.3, for any given s, we have:

$$\mathbb{E}[(v-s)_+ | a = a_i, i] = \Phi^c\left(\frac{s-a_i}{\tau_i}\right)(a_i - s) + \tau_i \phi\left(\frac{s-a_i}{\tau_i}\right). \tag{EC.99}$$

Hence, using $\phi'(x) = -x\phi(x)$ and the chain rule, we obtain:

$$\frac{\partial^- f(\boldsymbol{a})}{\partial a_m} = \Phi^c\left(\frac{s(\boldsymbol{a}) - a_m}{\tau_m}\right) + \underbrace{\frac{\partial^- s(\boldsymbol{a})}{\partial a_m}}_{\diamondsuit} \cdot \underbrace{\left[\Phi^c\left(\frac{s(\boldsymbol{a}) - a_w}{\tau_w}\right) - \Phi^c\left(\frac{s(\boldsymbol{a}) - a_m}{\tau_m}\right)\right]}_{\clubsuit}. \tag{EC.100}$$

From Claim EC.2, we have $\diamondsuit \leq 0$ for any $\boldsymbol{a}$. From Corollary 2- (a), we have $\clubsuit < 0$ at any $\boldsymbol{a}$ such that $f(\boldsymbol{a}) = 0$. Hence, inequality (EC.98) holds. We now observe that that, fixing a_w,

$$\lim_{a_m \to -\infty} f(\boldsymbol{a}) = -\mathbb{E}[(v - s(\boldsymbol{a}))_+ | a = a_w, w] < 0, \tag{EC.101}$$

$$\lim_{a_m \to \infty} f(\boldsymbol{a}) = \infty. \tag{EC.102}$$

Thus, due to continuity of $f(\boldsymbol{a})$ (which follows from the contintity of $s(\boldsymbol{a})$ — see Claim EC.2) and inequality (EC.98), we must have $\beta(a_w) = 1$. □

We now prove the uniqueness in the other direction, which will imply the main result we desired. Recall that we defined $\omega(a_m) := \{a_w : f(\boldsymbol{a}) = 0\}$. Fixing a_m, we observe that

$$\lim_{a_w \to -\infty} f(\boldsymbol{a}) = \mathbb{E}[(v - s(\boldsymbol{a}))_+ | a = a_m, m] > 0 \tag{EC.103}$$

$$\lim_{a_w \to \infty} f(\boldsymbol{a}) = -\infty. \tag{EC.104}$$

Hence, due to the continuity of $f(\boldsymbol{a})$, we have $|\omega(a_m)| \geq 1$. Now define $\overline{\omega}(a_m) = \max \omega(a_m)$. Due to continuity of f, $\overline{\omega}(a_m)$ is a continuous function. Furthermore, from (EC.103) and (EC.104), $f(a_m, \cdot)$ must change its sign from positive to negative near $\overline{\omega}(a_m)$. That is,

$$\left.\frac{\partial^- f}{\partial a_w}\right|_{(a_m, \overline{\omega}(a_m))} < 0 \tag{EC.105}$$

By the implicit function theorem, we then we deduce that $\overline{\omega}(\cdot)$ is a strictly increasing function:

$$\left.\frac{\partial^- \overline{\omega}(a_m)}{\partial a_m}\right|_{(a_m, \overline{\omega}(a_m))} = -\left.\frac{\partial^- f / \partial a_m}{\partial^- f / \partial a_w}\right|_{(a_m, \overline{\omega}(a_m))} > 0 \tag{EC.106}$$

where numerator is positive by Claim EC.4 and the denominator is negative by (EC.105).

We now establish that $|\omega(a_m)| = 1|$ for all a_m. Suppose for a contradiction that $|\omega(a_m)| > 1$ for some a_m. Then, for such a_m, there exists $\underline{\omega} \in \omega(a_m)$ such that $\underline{\omega} < \overline{\omega}(a_m)$. Because $\overline{\omega}(\cdot)$ is a strictly increasing function, its inverse $\overline{\omega}^{-1}$ is well-defined and also an increasing function, implying that $\overline{\omega}^{-1}(\underline{\omega}) < a_m$. At the same time, by definition of $\overline{\omega}(\cdot)$, we have $f(a_m, \underline{\omega}) = 0$ and $f(\overline{\omega}^{-1}(\underline{\omega}), \underline{\omega}) = 0$, and therefore $|\beta(\underline{\omega})| \geq 2$. This is a contradiction to Claim EC.4. Hence, we conclude that $|\omega(a_m)| = 1$. Finally, from inequality (EC.106), function $\omega(a_m)$ is strictly increasing. □

Building on Lemma EC.5, we now establish comparative statics of the optimal interview threshold $\boldsymbol{a}^*$ in the interview capacity C, which will be the main building block of proving Proposition EC.4.

LEMMA EC.6. *The optimal interview threshold $\boldsymbol{a}^*$, as a function of C, satisfy the following:*

(a) $\boldsymbol{a}^*$ *is continuous in* C.

(b) a_m^* *and* a_w^* *strictly decreases in* C.

(c) $a_m^* - a_w^*$ *strictly increases in* C.

Proof of Lemma EC.6-(a). We will use the Berge's Maximum Theorem.

DEFINITION EC.7 **(Section M.H of Mas-Colell et al. (1995)).**

1. *A correspondence* $\mathcal{X}: \Theta \rightrightarrows X$ *is upper hemicontinuous if it has a closed graph and images of compact sets are bounded.*
2. *A correspondence* $\mathcal{X}: \Theta \rightrightarrows X$ *is lower hemicontinuous if, for any* $\theta \in \Theta$, $x \in \mathcal{X}(\theta)$, *and open set* U *containing* x, *there exists a neighborhood* V *of* θ *such that* $\mathcal{X}(\theta) \cap U \neq \emptyset$ *for all* $\theta \in V$.

FACT EC.7 **(Berge's Maximum Theorem; Theorem M.K.6 of Mas-Colell et al. (1995)).** *Let* X *and* Θ *be topological spaces,* $f: X \times \Theta \to \mathbb{R}$ *be a continuous function, and* $\mathcal{X}: \Theta \rightrightarrows X$ *be a compact-valued correspondences such that* $\mathcal{X}(\theta) \neq \emptyset$ *for all* $\theta \in \Theta$. *Let* $\mathcal{X}^*(\theta) = argmax\{f(x,\theta) : x \in \mathcal{X}(\theta)\}$ *denote the set of maximizers. Then, if* $\mathcal{X}$ *is an upper- and lower hemicontinuous corresopndence, then* $\mathcal{X}^*$ *is upper hemicontinuous.*

If an upper-hemicontinuous correspondence is single-valued, it is a continuous function (Theorem M.H.1. of Mas-Colell et al. (1995)). Hence, because $\boldsymbol{a}^*$ is unique for each value of C (Proposition 1), it is a continuous function in C as long as Fact EC.7 can be applied. To apply Fact EC.7, it is convenient to reformulate the firm's problem (Benchmark) by introducing variable transformation $y_i = H_i(a_{1,m})$ for each $i \in \{m, w\}$ (we remind that H_i is C.D.F. of score for group i). Based on this variable transformation, it is straightforward to reformulate (Benchmark) as follows:

$$\max_{\mathbf{y}=(y_m,y_w)} V_1(\mathbf{y}) = \sum_{i\in\{m,w\}} 0.5 \int_{H_i^{-1}(y_i)} \int_{s(\mathbf{y})}^{\infty} v g_i(v \mid H_i^{-1}(y))\, \mathrm{d}v\, \mathrm{d}y \tag{EC.107}$$

$$\text{s.t.} \quad \mathbf{y} \in \mathcal{X}_1(C) \cap \mathcal{X}_2(C), \tag{EC.108}$$

where

$$s(\mathbf{y}) := \min\left\{ s \geq 0 : \sum_{i\in\{m,w\}} 0.5 \int_{H_i^{-1}(y_i)} \int_{s}^{\infty} g_i(v \mid H_i^{-1}(y))\, \mathrm{d}v\, \mathrm{d}y \leq \Delta \right\}, \tag{EC.109}$$

and $\mathcal{X}_1: (0,1) \rightrightarrows \mathbb{R}_+^2$ and $\mathcal{X}_2: (0,1) \rightrightarrows \mathbb{R}_+^2$ are correspondences defined as:

$$\begin{aligned} \mathcal{X}_1(C) &:= \{\mathbf{y} : 1 - 0.5y_m - 0.5y_w \leq C\}, \\ \mathcal{X}_2(C) &:= \{\mathbf{y} : \mathbf{y} \in [0,1]^2\}. \end{aligned} \tag{EC.110}$$

Specifically, $\mathcal{X}_1(C)$ is the interview capacity constraint.[20] The objective function (EC.107) is continuous in $\mathbf{y}$, due to continuity of $s(\mathbf{y})$ from Claim EC.2. Hence, it suffices to show that $\mathcal{X}_1(C) \cap \mathcal{X}_2(C)$ is a upper- and lower hemicontinuous correspondence in C. Because $\mathcal{X}_1(C) \cap \mathcal{X}_2(C)$ is closed and compact-valued for each C, its upper hemicontinuity is straightforward to deduce from Definition EC.7. For lower hemicontinuity, we will use the following.

[20] For the sake of applying Fact EC.7, it is more convenient to express the interview capacity constraint as inequality. Such inequality constraint is without loss of optimality because the interview constraint must bind at optimum.

FACT EC.8 **(From Exercise 3.12.d. of Stokey and Lucas Jr (1989))**. *Let $\phi : X \to Y$ and $\psi : X \to Y$ be lower-hemicontinuous correspondences, and define correspondence $\Gamma : X \to Y$ as $\Gamma(x) = \phi(x) \cap \psi(y)$. Suppose $\Gamma(x) \neq \emptyset$ and int $\phi(x) \cap$ int $\psi(x) \neq \emptyset$ for all $x \in X$.*[21] *Then, $\Gamma(\cdot)$ is lower hemicontinuous if ϕ and ψ are convex-valued.*

To apply the above fact, we first observe that both $\mathcal{X}_1(C)$ and $\mathcal{X}_2(C)$ are convex-valued (and nonempty for each C). Furthermore, $\mathcal{X}_2(C)$ is trivially lower-hemicontinuous in C as it does not depend on C. For lower-hemicontinuity of $\mathcal{X}_1(C)$, we note that $\mathcal{X}_1(C)$ is a half space, so it suffices to show that any half-space is lower-hemicontinuous, which we prove in the following:

FACT EC.9. *For $a \neq 0$, correspondence $S(b) := \{x \in \mathbb{R}^n : a \cdot x \leq b\}$ is lower-hemicontinuous in $b \in \mathbb{R}$. That is, for any $b_0 \in \mathbb{R}$, $x_0 \in S(b_0)$, and any open set $U \in \mathbb{R}^n$ containing x_0, there exists a neighborhood V of b_0 such that $S(b) \cap U \neq \emptyset$ for all $b \in V$.*

Proof of Fact EC.9. Without loss, let $U = \{x : ||x - x_0|| < \epsilon\}$ for some $\epsilon > 0$. Let $\delta = \epsilon ||a||$ and define $V = \{b : |b - b_0| < \delta\}$. We claim that, for any $b \in V$, there exists x' such that $x' \in S(b) \cap U$. To prove this, fix any $b \in V$ and define $x' := x_0 - \frac{b_0 - b}{||a||^2} a$. Then $||x' - x_0|| < \epsilon$ by definition of δ. Hence, $x' \in U$. Further, we observe that $a \cdot x' = a \cdot x_0 - (b_0 - b) \leq b$ because $x_0 \in S(b_0)$. Hence, $x' \in S(b)$ as well. □

Hence, by Fact EC.8, we deduce that $\mathcal{X}_1(C) \cap \mathcal{X}_2(C)$ is lower-hemicontinuous in C. Having proved that $\mathcal{X}_1(C) \cap \mathcal{X}_2(C)$ is upper- and lower hemicontinuous, we now apply Fact EC.7 and conclude that the optimal solution $\mathbf{y}^*$ is continuous in C. Because the optimal interview threshold is given by $a_i^* = H_i^{-1}(y_i^*)$ and $H_i(\cdot)$ is continuous function, it follows that $\boldsymbol{a}^*$ is continuous in C as well. □

Proof of Lemma EC.6-(b). By Proposition 1 and the definition of $\omega(a_m)$ in Lemma EC.5-(b), the optimal interview threshold $\boldsymbol{a}^* = (a_m^*, a_w^*)$ is the unique solution of the following:

$$\omega(a_m) = a_w, \tag{EC.111}$$

$$0.5(1 - H_m(a_m)) + 0.5(1 - H_w(a_w)) = C. \tag{EC.112}$$

We can further combine the above two equations to solve for a_m as follows:

$$\omega(a_m) = H_w^{-1}(2(1 - C) - H_m(a_m)). \tag{EC.113}$$

The right-hand side of the above equation is strictly decreasing in a_m and C. On the other hand, from Lemma EC.5-(b), $\omega(a_m)$ is a continuous and strictly increasing function (note that $\omega(\cdot)$ does not depend on C). Hence, by the standard monotone comparative statics argument, we deduce that a_m^* strictly decreases in C. Furthermore, because $a_w^* = \omega(a_m^*)$ and $\omega(\cdot)$ is strictly increasing, a_w^* strictly decreases in C as well. This completes the proof. □

[21] We use int(A) to denote the interior of set A.

Proof of Lemma EC.6-(c). Because a_i^* is continuous in C from Lemma EC.6-(a), it suffices to investigate a sign of the left derivative of $a_m^* - a_w^*$ with respect to C. Let $\frac{\mathrm{d}^- a_i^*}{\mathrm{d}C}$ denote the left derivative of a_i^* with respect to C. From Lemma EC.5, we recall that the optimal interview threshold $\boldsymbol{a}^*$ is the unique solution of the following equations:

$$f(\boldsymbol{a}) = 0,$$
$$0.5(1 - H_m(a_m)) + 0.5(1 - H_w(a_w)) = C.$$

Because the left partial derivative of f exists (Claim EC.2-(iii)), we can take the left derivative of the above two equations with respect to C, and solve for $\frac{\mathrm{d}^- a_i^*}{\mathrm{d}C}$. After some algebra, we obtain:

$$\frac{\mathrm{d}^- a_m^*}{\mathrm{d}C} = \frac{-2B}{AQ - BP}, \quad \frac{\mathrm{d}^- a_w^*}{\mathrm{d}C} = \frac{2A}{AQ - BP} \tag{EC.114}$$

where

$$A := \frac{\partial^- f(\boldsymbol{a})}{\partial a_m}\Big|_{\boldsymbol{a}=\boldsymbol{a}^*}, \quad B := \frac{\partial^- f(\boldsymbol{a})}{\partial a_w}\Big|_{\boldsymbol{a}=\boldsymbol{a}^*}, \quad P = -h_m(a_m^*), \quad Q = -h_w(a_w^*). \tag{EC.115}$$

In Claim EC.4, we have proved that $A > 0$. Furthermore, from part Lemma EC.6-(b), we have $\frac{\mathrm{d}a_w^*}{\mathrm{d}C} < 0$. Combining these observations and (EC.114), we deduce that $AQ - BP < 0$. Thus, to show that $a_m^* - a_w^*$ increases in C, it suffices to show that $A + B > 0$. To do so, by Fact EC.3 and straightforward algebras, we obtain

$$A = \Phi^c\left(\frac{s(\boldsymbol{a}^*) - a_m^*}{\tau_m}\right) + \left[\Phi^c\left(\frac{s(\boldsymbol{a}^*) - a_w^*}{\tau_w}\right) - \Phi^c\left(\frac{s(\boldsymbol{a}^*) - a_m^*}{\tau_m}\right)\right] \cdot \frac{\partial^- s(\boldsymbol{a})}{\partial a_m}\Big|_{\boldsymbol{a}=\boldsymbol{a}^*}, \tag{EC.116}$$

$$B = -\Phi^c\left(\frac{s(\boldsymbol{a}^*) - a_w^*}{\tau_w}\right) + \left[\Phi^c\left(\frac{s(\boldsymbol{a}^*) - a_w^*}{\tau_w}\right) - \Phi^c\left(\frac{s(\boldsymbol{a}^*) - a_m^*}{\tau_m}\right)\right] \cdot \frac{\partial^- s(\boldsymbol{a})}{\partial a_w}\Big|_{\boldsymbol{a}=\boldsymbol{a}^*}. \tag{EC.117}$$

Therefore,

$$A + B = \underbrace{\left[\Phi^c\left(\frac{s(\boldsymbol{a}^*) - a_m^*}{\tau_m}\right) - \Phi^c\left(\frac{s(\boldsymbol{a}^*) - a_w^*}{\tau_w}\right)\right]}_{\text{Term (1)}} \cdot \underbrace{\left[1 - \frac{\partial^- s(\boldsymbol{a})}{\partial a_m}\Big|_{\boldsymbol{a}=\boldsymbol{a}^*} - \frac{\partial^- s(\boldsymbol{a})}{\partial a_w}\Big|_{\boldsymbol{a}=\boldsymbol{a}^*}\right]}_{\text{Term (2)}}. \tag{EC.118}$$

Term (1) is positive by Lemma EC.5-(a), and term (2) is positive by Claim EC.2-(iv). Hence, we conclude that $A + B > 0$. This completes the proof. □

As the last building block toward proving Proposition EC.4, we show that, for a small enough Δ and $C = \Delta$, group m is always interviewed more than w.

Lemma EC.7. *There exists $\overline{\Delta}$ such that, for any $\Delta \le \overline{\Delta}$ and $C = \Delta$, $\boldsymbol{a}^* \in \mathsf{epi}(\gamma)$ where $\gamma(\cdot)$ is iso-interview curve (see Proposition EC.3).*

Proof of Lemma EC.7. To make the dependence of $\boldsymbol{a}^*$ on C and Δ clear, let $\boldsymbol{a}^*(C, \Delta)$ denote the Firm 1's optimal interview threshold given (C, Δ). Whenever $C = \Delta$, for $\boldsymbol{a}$ that satisfies the interview capacity constraint (2), we must have $s(\boldsymbol{a}) = 0$ because a mass of interviewed candidates with non-negative v is at most $C = \Delta$. Thus, the equation $f(\boldsymbol{a}) = 0$ in Lemma EC.5 reduces to:

$$\mathbb{E}[(v)_+ \mid a = a_m, m] = \mathbb{E}[(v)_+ \mid a = a_w, w] \tag{EC.119}$$

Using Fact EC.3, this equation is equivalent to

$$\Phi^c\left(-\frac{a_m}{\tau_m}\right)a_m+\tau_m\phi\left(-\frac{a_m}{\tau_m}\right)=\Phi^c\left(-\frac{a_w}{\tau_w}\right)a_w+\tau_w\phi\left(-\frac{a_w}{\tau_w}\right), \tag{EC.120}$$

and after some algebra, we deduce from the above equation that:

$$\frac{a_m}{a_w}=\frac{\Phi^c\left(-\frac{a_w}{\tau_w}\right)}{\Phi^c\left(-\frac{a_m}{\tau_m}\right)+\frac{\tau_m}{a_m}\phi\left(-\frac{a_m}{\tau_m}\right)-\frac{\tau_w}{a_m}\phi\left(-\frac{a_w}{\tau_w}\right)} \tag{EC.121}$$

From Proposition 1, the optimal interview threshold at $C=\Delta$ must satisfy eq. (EC.121). Furthermore, we observe that, as $C=\Delta\to 0$, any $\boldsymbol{a}$ that satisfy the interview constraint (2) increases to infinity. Hence, we have

$$\lim_{\Delta\to 0}\frac{a_m^*(\Delta,\Delta)}{a_w^*(\Delta,\Delta)}=1. \tag{EC.122}$$

Therefore, there exists $\overline{\Delta}$, such that, for all $\Delta\le\overline{\Delta}$, we have $\boldsymbol{a}^*(\Delta,\Delta)>0$ and $\frac{a_w^*(\Delta,\Delta)}{a_m^*(\Delta,\Delta)}>\frac{\sigma_w}{\sigma_m}$ (note that $\frac{\sigma_w}{\sigma_m}<1$ due to Assumption 1). Because $\gamma(a_m)=\frac{\sigma_w}{\sigma_m}a_w$ (Appendix EC.3.8), this implies that $\boldsymbol{a}^*(\Delta,\Delta)\in\mathsf{epi}(\gamma)$ for all $\Delta\le\overline{\Delta}$. This completes the proof. □

We are now ready to prove Proposition EC.4.

Proof of Proposition EC.4. We first show that there exists a value $\underline{C}$ such that $\eta_m(\boldsymbol{a}^*)=\eta_w(\boldsymbol{a}^*)$ if $C=\underline{C}$. By Lemma EC.7, for sufficiently small Δ, we have $\eta_m(\boldsymbol{a}^*)>\eta_w(\boldsymbol{a}^*)$ when $C=\Delta$. Moreover, since $a_m^*>a_w^*$ for all C (Corollary 1), it follows that $a_m^*>0>a_w^*$ at $C=0.5$. Because the iso-interview line $\gamma(a)=\frac{\sigma_w}{\sigma_m}a_m$ lies in the first or third quadrants, this implies that $\eta_m(\boldsymbol{a}^*)<\eta_w(\boldsymbol{a}^*)$ when $C=0.5$. By the continuity of $\boldsymbol{a}^*$ (Lemma EC.6-(a)) and the intermediate value theorem, there exists interview threshold $\underline{C}\in(\Delta,0.5)$ such that $\eta_m(\boldsymbol{a}^*)=\eta_w(\boldsymbol{a}^*)$ when $C=\underline{C}$.

We now show that such $\underline{C}$ is unique. Suppose, for contradiction, that there exist two distinct values of interview capacity, $\underline{C}$ and $\tilde{C}$ ($\underline{C}<\tilde{C}$), such that the optimal interview thresholds lie on the iso-interview curve for all $C\in\{\underline{C},\tilde{C}\}$. Let $\boldsymbol{a}^*=(a_m^*,a_w^*)$ and $\tilde{\boldsymbol{a}}=(\tilde{a}_m,\tilde{a}_w)$ denote the optimal interview thresholds at $C=\underline{C}$ and $C=\tilde{C}$, respectively. From Proposition EC.3 (specifically, the characterization of the iso-interview line $\gamma(a_m)=\frac{\sigma_w}{\sigma_m}a_m$), we deduce:

$$\frac{a_w^*}{a_m^*}=\frac{\tilde{a}_w}{\tilde{a}_m}=\frac{\sigma_w}{\sigma_m}.$$

Thus, we can express:

$$a_m^*-a_w^*=a_m^*\left(1-\frac{\sigma_w}{\sigma_m}\right),\quad \tilde{a}_m-\tilde{a}_w=\tilde{a}_m\left(1-\frac{\sigma_w}{\sigma_m}\right). \tag{EC.123}$$

On the other hand, Lemma EC.6-(b) establishes that the optimal interview thresholds for both groups strictly decrease in C. Thus, $a_m^*>\tilde{a}_m$. Combined with $\sigma_w<\sigma_m$ and the above equations, this implies that $a_m^*-a_w^*>\tilde{a}_m-\tilde{a}_w$. However, this contradicts Lemma EC.6-(c), which states that the gap $a_m^*-a_w^*$ must increase in C. Thus, the interview capacity threshold $\underline{C}$ must be unique.

Finally, as we argued earlier, $\eta_m(\boldsymbol{a}^*)-\eta_w(\boldsymbol{a}^*)$ changes its sign from positive to negative at $C=\underline{C}$. Thus, for any $C<1$, it holds that $\eta_m(\boldsymbol{a}^*)-\eta_w(\boldsymbol{a}^*)>0$ if and only if $C<\underline{C}$. This completes the proof. □

EC.3.10. Proof of Lemma 1

Recall from Lemma EC.5 that for any $\boldsymbol{a} = (a_m, a_w)$, we define the function $f(\boldsymbol{a}) := \mathbb{E}[(v - s(\boldsymbol{a}))_+ \mid a_m, m] - \mathbb{E}[(v - s(\boldsymbol{a}))_+ \mid a_m, w]$ and let $\omega(a_m) := \{a_w : f(\boldsymbol{a}) = 0\}$. In Lemma EC.5-(b), we showed that $\omega(a_m)$ is a strictly increasing function. Furthermore, by Lemma EC.2, we have $\omega(a_m^*) = a_w^*$, and by Corollary 2, it holds that $a_m^\rho \geq a_m^*$ and $a_w^* \geq a_w^\rho$. Thus, we must have $a_w^\rho \leq \omega(a_m^\rho)$. That is, (a_m^ρ, a_w^ρ) lies in the hypograph of $\omega(\cdot)$. Consequently, we obtain the inequality $\mathbb{E}[(v - s^\rho)_+ \mid a = a_m^\rho, m] \geq \mathbb{E}[(v - s^\rho)_+ \mid a = a_w^\rho, w]$. By mirroring the arguments presented in Lemma EC.5-(a), any $\boldsymbol{a}^\rho$ that satisfies this inequality must satisfy $\bar{G}_m(s^\rho \mid a_m^\rho) > \bar{G}_w(s^\rho \mid a_w^\rho)$. This completes the proof.

EC.3.11. Proof of Proposition 3

The implication for borderline candidates is straightforward. To prove that the strong candidate's hiring probability increases, we show that $s^\rho \leq s^*$ for any $\rho > \rho^*$. Assume that $s^\rho > 0$ since otherwise the claim is trivial. In the following, we establish that s^ρ strictly decreases in $\rho \geq \rho^*$. Define:

$$E(a_m, a_w, s) := \sum_{i \in \{m,w\}} 0.5 \int_{a_i}^{\infty} \bar{G}_i(s \mid a) h_i(a) \, \mathrm{d}a.$$

From Claim EC.2-(i) (Appendix EC.3.1), whenver $s^\rho > 0$, we have

$$E(a_m^\rho, a_w^\rho, s^\rho) = \Delta. \tag{EC.124}$$

By Corollary 2, for $\rho \geq \rho^*$, we have $a_w^\rho = H_w^{-1}(1 - 2\rho C)$ and $a_m^\rho = H_m^{-1}(1 - 2C(1-\rho))$, which implies

$$h_m(a_m^\rho) \cdot \frac{\mathrm{d}a_m^\rho}{\mathrm{d}\rho} = -h_w(a_w^\rho) \cdot \frac{\mathrm{d}a_w^\rho}{\mathrm{d}\rho} = 2C. \tag{EC.125}$$

Differentiating (EC.124) with respect to ρ and using (EC.125), we obtain, for any $\rho \geq \rho^*$,

$$\frac{\mathrm{d}s^\rho}{\mathrm{d}\rho} = \frac{2C\left[\bar{G}_w(s^\rho \mid a_w^\rho) - \bar{G}_m(s^\rho \mid a_m^\rho)\right]}{\int_{a_w^\rho}^{\infty} g_w(s^\rho \mid a) h_w(a) \, \mathrm{d}a + \int_{a_m^\rho}^{\infty} g_m(s^\rho \mid a) h_m(a) \, \mathrm{d}a}. \tag{EC.126}$$

From Lemma 1, the numerator in (EC.126) is negative. Thus, s^ρ strictly decreases for $\rho \geq \rho^*$ whenever $s^\rho > 0$. Further, observe that $s^\rho = s^*$ for $\rho \leq \rho^*$. Thus, we conclude that $s^\rho \leq s^*$ for $\rho \geq \rho^*$. Finally, the inequality must be strict whenever $s^* > 0$ because $\frac{\mathrm{d}s^\rho}{\mathrm{d}\rho}\Big|_{\rho=\rho^*} < 0$ by (EC.126) if $s^* > 0$.

EC.3.12. Group Unaware vs. Aware Benchmark

In our paper, we primarily focus on the group-aware approach (i.e., allowing the interview set to differ across groups) as our main benchmark. In this section, we show that a *group-unaware* approach—where the same interview set applies to both groups—is dominated by the group-aware approach in terms of *both* hiring diversity and the firm's objective.

Corollary EC.1 **(Comparison with Group-Unaware Strategy)**. *The group-aware solution* (Benchmark) *achieves a higher objective than the group-unaware one. Furthermore, the hiring mass of group w under the group-aware benchmark is always larger than that under the group-unaware approach.*

Proof of Corollary EC.1. The fact that the optimal group-aware solution has a higher objective value than the group-unaware one is straightforward since being group-unaware means that we add an extra constraint $a_m = a_w$ to (Benchmark). Let $\hat{a}$ denote the optimal interview threshold in the group-unaware benchmark.[22] Note that $\hat{a}$ is trivially given by the interview capacity constraint (2). However, by Corollary 1, we must have that $a_w^* < \hat{a} < a_m^*$. By Lemma EC.4, this directly implies that group w's hiring mass under the group-unaware solution must be strictly less than that under the group-aware solution. □

EC.4. Proofs Related to Section 4

EC.4.1. Proof of Proposition 4

Similar to Firm 1's case, we apply Propositions EC.1 and EC.2 by setting $\Psi_i(a)$ in (OPT-Meta) as $\Psi(a \mid \boldsymbol{a}_1^*, s_1^*)$ (defined in (13)). Let $F_{2,i}(a, s_2) = \mathbb{E}[(v - s_2)_+ \mid a, i]\Psi_i(a \mid \boldsymbol{a}_1^*, s_1^*)$ denote the discounted excess value function for a given s_2. Since the function $\Psi(a \mid \boldsymbol{a}_1^*, s_1^*)$ has only one jump discontinuity at $a = a_{1,i}^*$, Assumption EC.1-(ii) is satisfied. We now verify Assumption EC.1-(iii). By Fact EC.2, for any s_2, the function $F_{2,i}(a, s_2)$ is increasing for $a < a_{1,i}^*$. In what follows, we show that $F_{2,i}(a, s_2)$ is unimodal for $a > a_{1,i}^*$.

Claim EC.5. *Let $v \mid a \sim \mathcal{N}(a, \tau^2)$ and denote its p.d.f. and C.D.F. by $g(\cdot|a)$ and $G(\cdot|a)$, respectively. Then, for any s_1 and s_2, function $f(a) := \mathbb{E}[(v - s_2)_+ \mid a]G(s_1|a)$ is unimodal in $a \in \mathbb{R}$.*

Proof of Claim EC.5. From Fact EC.3, we have $f'(a) = (1 - G(s_2 \mid a))G(s_1|a) - \mathbb{E}[(v - s_2)_+ \mid a]g(s_1|a)$. Thus,

$$f'(a) = \bar{G}(s_2 \mid a)g(s_1 \mid a)\underbrace{\left(\frac{G(s_1 \mid a)}{g(s_1 \mid a)} - \mathbb{E}[v - s_2 \mid v - s_2 \geq 0, a]\right)}_{:=D(a)} \tag{EC.127}$$

By asymptotic properties of the Mills ratio (Sampford 1953), we have $\lim_{a\to-\infty} D(a) > 0$ and $\lim_{a\to\infty} D(a) = 0$. Thus, the intermediate value theorem implies that there must exists a critical point a such that $f'(a) = 0$. We claim that such critical point must be unique. To show this, note that $f'(a) = 0$ if and only if

$$\mathbb{E}[v - s_2 \mid v - s_2 \geq 0, a] = \frac{G(s_1 \mid a)}{g(s_1 \mid a)}. \tag{EC.128}$$

We claim that the left- and right-hand sides are increasing and decreasing functions in a, respectively, which in turn implies that the function f must be uniomodal. To show that the left hand side is increasing in a, we observe that the family of normal distributions $\mathcal{N}(a - s_2, \tau^2)$ is increasing in a in the likelihood ratio order from Fact EC.2. Therefore, by Theorem 1.C.6 of Shaked and Shanthikumar (2007), $\mathbb{E}[v - s_2 \mid v - s_2 \in A, a]$ increases in a for any set A, and so does $\mathbb{E}[v - s_2 \mid v - s_2 \geq 0, a]$.

[22] It is straightforward to show that the optimal interview set under the group-unaware approach must be greedy in the score, by mirroring arguments of Proposition EC.1.

For the right-hand side, let $R(x)=\phi(x)/\Phi(x)$ denote the inverse Mills ratio of the standard normal distribution. Then the RHS is decreasing in a if and only if $1/R(\frac{s_1-a}{\tau})$ is. The inverse Mills ratio $R(x)$ is decreasing in x (Sampford 1953). Hence, $1/R(\frac{s_1-a}{\tau})$ is decreasing in a, implying that the RHS is a decreasing function in a. □

Thus, applying Propositions EC.1 and EC.2, we conclude that Firm 2's optimal interview set $A_2^*=(A_{2,m},A_{2,w}^*)$ and optimal hiring threshold s_2^* uniquely exist and satisfy the optimality conditions (EC.4) and (EC.5). By (EC.4), $A_{2,w}^*$ is a superlevel set of $F_{2,i}(a,s_2^*)$ with a common level θ, i.e., $A_{2,i}^*=\{a: F_{2,i}(a,s_2^*)\geq\theta\}$, where the level θ is uniquely determined by the constraint (15). Furthermore, since $F_{2,i}(a,s_2^*)$ is increasing for $a<a_{1,i}^*$ and unimodal for $a>a_{1,i}^*$, the superlevel set must consist of a greedy interval below $a_{1,i}^*$ and a non-greedy interval above $a_{1,i}^*$. Equivalently, A_{2,i^*} must be of the form $A_{2,i}^*=[b_{2,i}^*,a_{1,i}^*]\cup[c_{2,i}^*,d_{2,i}^*]$, where $-\infty<b_{2,i}^*\leq a_{1,i}^*\leq c_{2,i}^*\leq d_{2,i}^*<\infty$. This proves part (a). Part (b) follows directly from the optimality condition (EC.5). This completes the proof.

EC.4.2. Proof of Proposition 5

First, as we argued in Section 4.3, we observe that if $C_f=\Delta_f$, then firm f must set $s_f^*=0$ and $s_f^\rho=0$ whenever $C_f=\Delta_f$. Given this observation, we proceed in two steps. First, we show that, for small enough $C_f=\Delta_f$, Firm 2 never competes with Firm 1 under (Benchmark). In the second step, we show that, under (Intervention), Firm 2 replaces all borderline group w candidates with the borderline group m candidates in their optimal interview set. As such, because $s_f^*=s_f^\rho=0$, we must have $\lambda_{2,w}^*>\lambda_{2,w}^\rho$ for all $\rho>\rho^*$. We elaborate each step in the following.

Step 1: For small enough $C_f=\Delta_f$, Firm 2 does not compete with Firm 1. For notational brevity, define

$$F_{2,i}(a):=\mathbb{E}[(v)_+\mid a,i]\Psi_i(a\mid \boldsymbol{a}_1^*,0). \tag{EC.129}$$

By Proposition 4, Firm 2's optimal interview set under (Benchmark) is given by

$$A_{2,i}^*=\{a: F_i(a)\geq\theta(C_2)\}, \tag{EC.130}$$

where $\theta(C_2)$ is uniquely determined such that the total mass of the interview set is C_2, i.e.,

$$\sum_{i\in\{m,w\}}0.5\int_{\{a:F_{2,i}(a)\geq\theta(C_2)\}}dH_i(a)=C_2. \tag{EC.131}$$

We now specify φ_1 and φ_2 such that if $C_f=\Delta_f\leq\varphi_f$, Firm 2 does not compete with Firm 1. To begin, define

$$k_i:=\max_{a\in\mathbb{R}}\mathbb{E}[(v)_+\mid a,i]G_i(0\mid a). \tag{EC.132}$$

Note that k_i is independent of the capacity parameters. Further, because $a_{1,i}^*$ decreases in C_1 (Lemma EC.6-(b)), there exists φ_1 such that for all $C_1\leq\varphi_1$, we have $k_i\leq\lim_{a\uparrow a_{1,i}^*}F_i(a)$. We now define φ_2. Note that $\theta(C_2)$ in equation (EC.130) decreases in C_2. Hence, there exists φ_2 such that for all $C_2\leq$

φ_2, it holds that $\theta(C_2) \geq \max\{k_m, k_w\}$. Putting everything together, it follows that $A^*_{2,i} \cap [a^*_{1,i}, \infty) = \emptyset$ for all $i \in \{m, w\}$ whenever $C_f = \Delta_f \leq \varphi_f$ for all $f \in \{1, 2\}$ (see Figure 4 in Section 4.3 for illustration).

Step 2: Firm 2 hires fewer group w under (Intervention) for all $\rho > \rho^*$. Similar to $F_{2,i}(a)$ in the previous step, for $\rho > \rho^*$, define a function:

$$F_{2,i}(a \mid \rho) = \mathbb{E}[(v)_+ \mid a, i]\Psi_i(a \mid \boldsymbol{a}^*_1(\rho), 0). \tag{EC.133}$$

Then, by Proposition 4, Firm 2's optimal interview set $A^*_{2,i}(\rho)$ under (Intervention) is given by

$$A^\rho_{2,i} = \{a : F_i(a \mid \rho) \geq \theta(C_2 \mid \rho)\}, \tag{EC.134}$$

where $\theta(C_2 \mid \rho)$ is uniquely determined such that the total mass of the interview set is C_2, i.e.,

$$\sum_{i \in \{m,w\}} 0.5 \int_{\{a : F_{2,i}(a|\rho) \geq \theta(C_2|\rho)\}} dH_i(a) = C_2. \tag{EC.135}$$

We now claim that $\theta(C_2) = \theta(C_2 \mid \rho)$. To see why, we recall that (i) $a^\rho_{1,m}$ ($a^\rho_{1,w}$, resp.) strictly increases (decreases, resp.) in $\rho \geq \rho^*$ (Corollary 2) and (ii) $F_i(a \mid \rho)$ increases in $a < a^\rho_{1,i}$. Thus, we have

$$\inf_{a \in (a^*_{1,m}, a^\rho_{1,m})} F_m(a \mid \rho) = F_m(a^*_{1,m} \mid \rho) > \theta(C_2). \tag{EC.136}$$

On the other hand, because $F_{2,w}(a \mid \rho)$ is unimodal on $a > a^\rho_{1,w}$, and due to the definition of k_i in (EC.132), we have

$$\theta(C_2) \geq k_i > \sup_{a \in (a^\rho_{1,w}, a^*_{1,w})} F_w(a \mid \rho) \tag{EC.137}$$

for all $\rho > \rho^*$ (see Figure 4 in Section 4.3 for illustration).

Finally, we note that $F_{2,i}(a \mid \rho) = F_{2,i}(a)$ except for all a of the borderline candidates (i.e. $a \in [a^*_{1,m}, a^\rho_{1,m}]$ for group m and $a \in [a^\rho_{1,w}, a^*_{1,w}]$ for group w) because $s^*_f = s^\rho_f = 0$, $f \in \{1, 2\}$. Combined with (EC.136) and (EC.137), this implies that

$$\begin{aligned} \{a : F_m(a \mid \rho) \geq \theta(C_2)\} &= A^*_{2,m} \cup [a^*_m, a^\rho_m] \\ \{a : F_w(a \mid \rho) \geq \theta(C_2)\} &= A^*_{2,w} \setminus [a^\rho_w, a^*_w]. \end{aligned} \tag{EC.138}$$

However, we have:

$$\begin{aligned} \sum_{i \in \{m,w\}} 0.5 \int_{\{a : F_i(a|\rho) \geq \theta(C_2)\}} \mathrm{d}H_i(a) &= \sum_{i \in \{m,w\}} 0.5 \int_{A^*_{2,i}} \mathrm{d}H_i(a) + 0.5 \left(\int_{a^*_{1,m}}^{a^\rho_{1,m}} \mathrm{d}H_m(a) - \int_{a^\rho_{1,w}}^{a^*_{1,w}} \mathrm{d}H_w(a) \right) \\ &= C_2, \end{aligned}$$

where the equality is due to $\int_{a^*_{1,m}}^{a^\rho_{1,m}} \mathrm{d}H_m(a) = \int_{a^\rho_{1,w}}^{a^*_{1,w}} \mathrm{d}H_w(a)$ (by Corollary 2), and the total mass of $(A^*_{2,m}, A^*_{2,w})$ is C_2 by definition. Hence, by (EC.134) and (EC.135), we conclude that $\theta(C_2 \mid \rho) = \theta(C_2)$. As such, Firm 2's interview set under (Intervention) must be given by (EC.138). That is, Firm 2 replaces all of borderline group w with borderline group m. Combing with the fact that the hiring threshold for Firm 2 remains zero when $C_2 = \Delta_2$, it follows that Firm 2 hires strictly fewer group w candidates under (Intervention) for all $\rho > \rho^*$. This completes the proof.

EC.4.3. Proof of Proposition 6

Define the following set:

$$D_w^\rho := [a_{1,w}^\rho, a_{1,w}^*] \cap A_{2,w}^* \setminus A_{2,w}^\rho \tag{EC.139}$$

$$D_m^\rho := [a_{1,m}^*, a_{1,m}^\rho] \cap A_{2,m}^\rho \setminus A_{2,m}^* \tag{EC.140}$$

That is, D_w^ρ is borderline group w candidates who lose an interview spot from Firm 2 due to Firm 1's adoption of the ρ-Rooney rule. Similarly, D_m^ρ is borderline group m candidates who gain an interview spot from Firm 2. Let $\mu_i^\rho(C_1, C_2)$ denote a mass of D_i^ρ under the interview capacities (C_1, C_2) and the ρ-Rooney rule. In Proposition 5, we showed that $\mu_i^\rho(\Delta_1, \Delta_2) > 0$ for all $\rho > \rho^*$. We will extend this result for all (C_1, C_2) in a neighborhood of (Δ_1, Δ_2). To do so, in the following lemma, we first show that the function $\mu_i^\rho(C_1, C_2)$ is continuous in (C_1, C_2) and ρ. To show this, it suffices to establish that the endpoints of $A_{2,i}^*$ and $A_{2,i}^\rho$ (characterized by Proposition 4) vary continuously with (C_1, C_2) and ρ. We establish this in the following lemma, proven at the end of this section.

LEMMA EC.8. *The end points of $A_{2,i}^*$ and $A_{2,i}^\rho$ are continuous in (C_1, C_2) and ρ.*

We combine the above lemma with the following auxiliary claims to deduce our desired result. Recall that, we defined ρ^* as Firm 1's interview fraction of group w under the benchmark (Equation (4)). Fixing Δ_1, this interview fraction is a function of the interview capacity $C_1 \geq \Delta_1$. Hereafter, we write $\rho^*(C_1)$ to explicitly indicate its dependence on C_1 (for a fixed Δ_1). In the following claim, we show that this fraction increases when $C_1 = \Delta_1$ for sufficiently small Δ_1.

CLAIM EC.6. *For sufficiently small Δ_1, we have $\frac{\mathrm{d}\rho^*(C_1)}{\mathrm{d}C_1}\Big|_{C_1=\Delta_1} > 0$.*

The next claim allows us to use the continuity of $\mu_i^\rho(C_1, C_2)$ to extend the positivity of $\mu_i^\rho(\Delta_1, \Delta_2)$ to a broader range of (C_1, C_2) *whenever the ρ-Rooney rule applies.*

CLAIM EC.7. *Let $f : \mathbb{R}^3 \to \mathbb{R}$ and $\rho^* : \mathbb{R} \to \mathbb{R}$ be continuous functions. Suppose there is a point $(\Delta_1, \Delta_2) \in \mathbb{R}^2$ such that (i) for every $\rho > \rho^*(\Delta_1)$, we have $f(\Delta_1, \Delta_2, \rho) > 0$ and (ii) $\frac{\mathrm{d}\rho^*(C_1)}{\mathrm{d}C_1}\Big|_{C_1=\Delta_1} > 0$. Then there exists $\bar{C}_f > \Delta_f$, $f \in \{1, 2\}$, such that for all $C_f \in [\Delta_f, \bar{C}_f]$, we have $f(C_1, C_2, \rho) > 0$ whenever $\rho > \rho^*(C_1)$.*

The proofs of Claims EC.6 and EC.7 can be found at the end of this section.

To complete the proof of Proposition 6, fix small enough (Δ_1, Δ_2). In the proof of Proposition 5, we showed that at $C_1 = \Delta_1$ and $C_2 = \Delta_2$, Firm 2 replaces all borderline group w candidates in its interview set with all borderline group m whenever Firm 1 adopts the ρ-Rooney rule with $\rho > \rho^*(\Delta_1)$. In other words, $\mu_i^\rho(\Delta_1, \Delta_2) > 0$ for any $\rho > \rho^*(\Delta_1)$ and all $i \in \{m, w\}$. Because Firm 1's interview threshold is continuous in C_1 (Lemma EC.6–(a)), it follows that $\rho^*(C_1)$ is continuous in C_1. Since $\mu_i^\rho(C_1, C_2)$ varies continuously in (C_1, C_2) and ρ (by Lemma EC.8) and $\rho^*(C_1)$ is strictly increasing at $C_1 = \Delta_1$ (by

Claim EC.6), we apply Claim EC.7 to conclude that there exists $\bar{C}_f > \Delta_f$ such that if $C_f \in [\Delta_f, \bar{C}_f]$, we have $\mu_i^\rho(C_1, C_2) > 0$ for all $\rho > \rho^*(C_1)$ and both $i \in \{m, w\}$.

Proof of Lemma EC.8. Similar to Firm 1's case (Lemma EC.6–(a)), the proof follows from an application of Berge's maximum theorem (Fact EC.7), although verifying the conditions for applying Fact EC.7 is more intricate. We focus on showing the continuity of the endpoints of $A_{2,i}^*$ in (C_1, C_2). The argument for $A_{2,i}^\rho$ and its continuity in ρ proceeds in a similar way.

From Proposition 4, Firm 2's optimal interview set is given by $A_{2,i}^* = [b_{2,i}^*, a_{1,i}^*] \cup [c_{2,i}^*, d_{2,i}^*]$ for some $b_{2,i}^* \le a_{1,i}^* \le c_{2,i}^* \le d_{2,i}^*$. As such, Firm 2's optimal interview set is characterized by a set of end points $\mathcal{E}^* = \{b_{2,i}^*, c_{2,i}^*, d_{2,i}^* : i \in \{m, w\}\}$. To establish continuity of $\mathcal{E}^*$ in (C_1, C_2), we first introduce the change of variables $y_{1,i} := H_i(b_{2,i})$, $y_{2,i} := H_i(c_{2,i})$, and $y_{3,i} := H_i(d_{2,i})$. Further, let $\mathbf{y}_i = \{y_{k,i} : k \in \{1,2,3\}, i \in \{m, w\}\}$ and $\mathbf{y} = (\mathbf{y}_m, \mathbf{y}_w) \in \mathbb{R}^6$. We can then express Firm 2's interview set as $A_{2,i}(\mathbf{y}) := [H_i^{-1}(y_{1,i}), a_{1,i}^*] \cup [H_i^{-1}(y_{2,i}), H_i^{-1}(y_{3,i})]$, and the Firm 2's optimization problem is equivalent to finding the optimal (transformed) end points $\mathbf{y}^* \in \mathbb{R}^6$. With a slight abuse of notation (and in a similar vein of Claim EC.2), we further define function $s_2(\cdot) : \mathbb{R}^6 \to \mathbb{R}$,

$$s_2(\mathbf{y}) = \min \Big\{ s \ge 0 : \sum_{i \in \{m,w\}} 0.5 \int_{A_{2,i}(\mathbf{y})} \int_s^\infty g_i(v \mid a) h_i(a) \Psi_i(a \mid \boldsymbol{a}_1^*, s_1^*) \, \mathrm{d}v \, \mathrm{d}a \le \Delta \Big\}. \tag{EC.141}$$

Then, the optimal transformed end point $\mathbf{y}^* = (\mathbf{y}_m^*, \mathbf{y}_w^*)$ solves the following optimization problem:

$$\max_{\substack{\mathbf{y}_i = (y_{1,i}, y_{2,i}, y_{3,i}) \\ i \in \{m,w\}}} \sum_{i \in \{m,w\}} 0.5 \int_{A_{2,i}(\mathbf{y})} \int_{s_2(\mathbf{y})}^\infty v g_i(v \mid H_i^{-1}(y)) \Psi_i(H_i^{-1}(y) \mid \boldsymbol{a}_1^*, s_1^*) \, \mathrm{d}v \, \mathrm{d}y \tag{EC.142}$$

$$\text{s.t.} \quad \mathbf{y} \in \mathcal{Y}_1(C_1, C_2) \cap \mathcal{Y}_2(C_1, C_2), \tag{EC.143}$$

where (i) we recall the definition of function $s_2(\cdot)$ from line (EC.141) and (ii) $\mathcal{Y}_1 : (0,1)^2 \rightrightarrows \mathbb{R}_+^6$ and $\mathcal{Y}_2 : (0,1)^2 \rightrightarrows \mathbb{R}_+^6$ are correspondences defined as:

$$\mathcal{Y}_1(C_1, C_2) = \{(\mathbf{y}_m, \mathbf{y}_w) \in [0,1]^6 : y_{1,i} \le H_i(a_{1,i}^*) \le y_{2,i} \le y_{3,i}, \ \forall i \in \{m, w\}\} \tag{EC.144}$$

$$\mathcal{Y}_2(C_1, C_2) = \Big\{ (\mathbf{y}_m, \mathbf{y}_w) \in [0,1]^6 : \sum_{i \in \{m,w\}} 0.5(H_i(a_{1,i}^*) - y_{1,i} + y_{3,i} - y_{2,i}) \le C_2 \Big\} \tag{EC.145}$$

Here, $\mathcal{Y}_1(C_1, C_2)$ enforces the condition $b_{2,i} \le a_{1,i}^* \le c_{2,i} \le d_{2,i}$ and $\mathcal{Y}_2(C_1, C_2)$ represents the interview capacity constraint (15).[23] We now apply Fact EC.7 to deduce the continuity of $\mathbf{y}^*$ in (C_1, C_2), which implies continuity of $\mathcal{E}^*$ by the continuity of H_i. Since $s_2(\mathbf{y})$ is continuous,[24] the objective function (EC.142) is continuous in $\mathbf{y}$. Next, we establish that the constraint set $\mathcal{Y}_1(C_1, C_2) \cap \mathcal{Y}_2(C_1, C_2)$ is upper- and lower-hemicontinuous (see Definition EC.7). Upper hemicontinuity is immediate because $a_{1,i}^*$ is continuous in C_1 (Lemma EC.6-(a)) and both $\mathcal{Y}_1(C_1, C_2)$ and $\mathcal{Y}_2(C_1, C_2)$ are compact-valued. For lower-hemicontinuity, note that $\mathcal{Y}_1(C_1, C_2)$ and $\mathcal{Y}_2(C_1, C_2)$ can each be expressed as the intersection of

[23] The inequality is without loss because, at optimum, the interview capacity constraint must bind.

[24] One can mirror the prof of Claim EC.2 to show that $s_2(\mathbf{y})$ is continuous in $\mathbf{y}$.

finitely many half-spaces of the form $\{\mathbf{y}:\mathbf{c}\cdot\mathbf{y}\le b(C_1,C_2)\}$, where $\mathbf{c}$ is a constant vector and $b(C_1,C_2)$ is a continuous function. By Fact EC.8,[25] it suffices to show that each of such half-spaces is lower-hemicontinuous.

Claim EC.8. *For $\mathbf{c}\in\mathbb{R}^n$ and $\mathbf{C}\in\mathbb{R}^m$, let $b(\mathbf{C})$ be a continuous function in $\mathbf{C}$ and define $S(\mathbf{C}):=\{\mathbf{x}\in\mathbb{R}^n:\mathbf{c}\cdot\mathbf{x}\le b(\mathbf{C})\}$. Then $S(\mathbf{C})$ is lower-hemicontinuous in $\mathbf{C}$. That is, for any $\mathbf{C}_0\in\mathbb{R}^m$, $\mathbf{x}_0\in S(\mathbf{C}_0)$, and any open set V containing $\mathbf{x}_0$, there exists a neighborhood U of $\mathbf{C}_0$ such that $S(\mathbf{C})\cap V\neq\emptyset$ for all $\mathbf{C}\in U$.*

Proof of Claim EC.8. Without loss of generality, let $V=\{\mathbf{y}\in\mathbb{R}^n:||\mathbf{y}-\mathbf{x}_0||<\epsilon\}$ for some $\epsilon>0$. For such ϵ, because $b(\mathbf{C})$ is a continuous function, there exists a neighborhood U of $\mathbf{C}_0$ such that $|b(\mathbf{C})-b(\mathbf{C}_0)|<\frac{\epsilon||\mathbf{c}||}{4}$ for all $\mathbf{C}\in U$. Take $\mathbf{y}':=\mathbf{x}_0-\frac{\epsilon}{2||\mathbf{c}||}\mathbf{c}$. We claim that $\mathbf{y}'\in S(\mathbf{C})\cap V$ whenever $\mathbf{C}\in U$. Clearly, $||\mathbf{y}'-\mathbf{x}_0||=\frac{\epsilon}{2}$, and thus $\mathbf{y}'\in V$. Further,

$$\mathbf{c}\cdot\mathbf{y}'=\mathbf{c}\cdot\mathbf{x}_0-\frac{\epsilon}{2}||\mathbf{c}||\le b(\mathbf{C}_0)-\frac{\epsilon}{2}||\mathbf{c}||\le b(\mathbf{C})-\frac{\epsilon}{4}||\mathbf{c}||. \tag{EC.146}$$

The first inequality follows because $\mathbf{x}_0\in S(\mathbf{C}_0)$. The last inequality follows from the definition of U. Hence, we deduce that $\mathbf{y}'\in S(\mathbf{C})$, as desired. □

Therefore, we can apply Fact EC.7 to the optimization problem (EC.142)-(EC.143) and deduce that the optimal solution $\mathbf{y}^*$ is continuous in (C_1,C_2). Hence, the end points of $A^*_{2,i}$ are continuous. This completes the proof. □

Proof of Claim EC.6. From (EC.122) in Appendix EC.3.9, $a^*_{1,m}/a^*_{1,w}\to 1$ if $C_1=\Delta_1$ and as $\Delta_1\to 0$. Thus, for small enough $C_1=\Delta_1$, we can approximate $a^*_{1,m}=a^*_{1,w}=a^*$ for some $a^*>0$. In this regime, we can express $\rho^*(C_1)=(R(a^*)+1)^{-1}$ where $R(a^*):=\frac{\Phi^c(a^*/\sigma_m)}{\Phi^c(a^*/\sigma_w)}$ (we remind that we use ϕ and Φ to denote the p.d.f. and C.D.F. of standard normal distribution, respectively). Note that a^* decreases in C_1 because $0.5\Phi^c(a^*/\sigma_m)+0.5\Phi^c(a^*/\sigma_w)=C_1$. Thus, to prove $\frac{\mathrm{d}\rho^*}{\mathrm{d}C_1}\Big|_{C_1=\Delta_1}>0$, it suffices to show that $R(a^*)$ increases in $a^*>0$. Let $h(x)=\phi(x)/\Phi^c(x)$ denote the hazard ratio function. The ratio $R(a^*)$ increases in a^* for $a^*>0$ if and only if $\frac{h(a^*/\sigma_w)}{h(a^*/\sigma_m)}>\frac{\sigma_w}{\sigma_m}$. Since $h(x)$ is an increasing function (Sampford 1953) and $\sigma_m>\sigma_w$, we conclude that $R(a^*)$ increases in $a^*>0$, as desired. □

Proof of Claim EC.7. Suppose, for contradiction, that no such $C_f>\Delta_f$ exists. Then for every $n\in\mathbb{N}$, we can pick $(C_{1,n},C_{2,n})$ and $\rho_n>\rho^*(C_{1,n})$ such that $C_{1,n}\in[\Delta_1,\Delta_1+1/n]$, $C_{2,n}\in[\Delta_2,\Delta_2+1/n]$, but $f(C_{1,n},C_{2,n},\rho_n)\le 0$. We will show that no such "bad" sequence exists using continuity of f and ρ^*. Let ρ_∞ be a limit of $\{\rho_n\}$, so $(C_{1,n},C_{2,n},\rho_n)\to(\Delta_1,\Delta_2,\rho_\infty)$. Since $\rho_n>\rho^*(C_{1,n})$ for all n, continuity of ρ^* implies $\rho_\infty\ge\rho^*(\Delta_1)$. We now consider the following two cases.

Case 1: $\rho_\infty>\rho^*(\Delta_1)$. By hypothesis (i), we have $f(\Delta_1,\Delta_2,\rho_\infty)>0$. However, by continuity of f, it holds that $f(C_{1,n},C_{2,n},\rho_n)\to f(\Delta_1,\Delta_2,\rho_\infty)>0$, contradicting $f(C_{1,n},C_{2,n},\rho_n)\le 0$.

[25] Because a finite intersection of convex sets is convex, Fact EC.8 implies that a finite intersection of convex-valued lower-hemicontinuous correspondences is also lower-hemicontinuous.

Case 2: $\rho_\infty = \rho^*(\Delta_1)$. By hypothesis (ii), we have $\rho^*(C_{1,n}) \geq \rho^*(\Delta_1)$ for large enough n. Since $\rho_n > \rho^*(C_{1,n})$ by construction, it follows that $\rho_n > \rho^*(\Delta_1)$ for large n. Thus, by hypothesis (i), $f(\Delta_1, \Delta_2, \rho_n) > 0$ for all large n. However, by continuity of f and $(C_{1,n}, C_{2,n}) \to (\Delta_1, \Delta_2)$ as $n \to \infty$, it holds that $f(C_{1,n}, C_{2,n}, \rho_n) > 0$ eventually, again contradicting $f(C_{1,n}, C_{2,n}, \rho_n) \leq 0$.

Thus, no such sequence $\{(C_{1,n}, C_{2,n}, \rho_n)\}$ exists, implying that there must exist $\bar{C}_f > \Delta_f$, $f \in \{1,2\}$ such that for all $C_f \in [\Delta_f, \bar{C}_f]$, we have $f(C_1, C_2, \rho) > 0$ whenever $\rho > \rho^*(C_1)$. □

EC.4.4. Proof of Proposition 7

Fix a value of C_1 and $\rho > \rho^*$. Part (a) follows immediately from $s_1^\rho \leq s_1^*$ as established in Proposition 3 (Appendix EC.4.4). To establish parts (b) and (c), we show that for a given C_1, there exists $\bar{C}_2$ such that if $C_2 \leq \bar{C}_2$,[26]

$$\max\{s_2^*, s_2^\rho\} \leq s_1^\rho. \tag{EC.147}$$

Note that (EC.147) implies parts (b) and (c). To see why, let D_w^ρ denote the subset of borderline group w candidates who lose an interview spot from Firm 2 due to Firm 1's adption of the ρ Rooney rule (see (EC.139)). By definition, such candidates gain an interview spot from Firm 1. Consequently, for any group w candidate with score $a \in D_w^\rho$, their hiring probability changes from $\bar{G}_w(s_2^* \mid a)$ to $\bar{G}_w(s_1^\rho \mid a)$. Thus, for $C_2 \leq \bar{C}_2$, their hiring probability decreases, implying part (b). Similarly, define D_m^ρ as the subset of borderline group m candidates who gain an interview spot from Firm 2. Noting that $\min\{s_1^\rho, s_1^*\} = s_1^\rho$ (Proposition 3), we deduce part (c) from inequality (EC.147).

To prove (EC.147), we show that Firm 2's optimal hiring threshold increases in C_2.

Lemma EC.9. *Both s_2^* and s_2^ρ increase in $C_2 \in [\Delta_2, 1]$.*

Proof of Lemma EC.9. We provide the proof for the case of s_2^*. The proof s_2^ρ follows a similar argument. Let $C_2 < \tilde{C}_2$. With a slight abuse of notation, we use s_2^* and $\tilde{s}_2$ to denote the optimal hiring threshold given C_2 and $\tilde{C}_2$, respectively. Given these notations, it suffices to show that $s_2^* \leq \tilde{s}_2$.

Toward that goal, we recall notation from Definition EC.1 and Definition EC.2 (Appendix EC.2.2). Specifically, letting $F_{2,i}(a, s_2) := \mathbb{E}[(v - s_2)_+ | a, i] \Psi_i(a | \boldsymbol{a}_1^*, s_1^*)$, set $B_i(s_2, C_2)$ from Definition EC.1 is given by:

$$B_i(s_2, C_2) := \{a \in \mathbb{R} : F_{2,i}(a, s_2) \geq \theta(C_2)\}, \quad \forall i \in \{m, w\}. \tag{EC.148}$$

where $\theta(C_2)$ is the unique solution of the following equation in θ (and parameterized by s_2 and C_2):

$$\sum_{i \in \{m,w\}} 0.5 \int_{\{a : F_{2,i}(a, s_2) \geq \theta\}} h_i(a)\, \mathrm{d}a = C_2, \tag{EC.149}$$

[26] Note that we fix a value of C_1 and consider a fixed $\rho > \rho^*$, where ρ^* depends on C_1. Thus, the interview capacity threshold $\bar{C}_2$ may depend on C_1 and ρ.

Furthermore, from Definition EC.2, we recall the following function $\lambda(B,s)$ given interview set $B=(B_m,B_w)$ and (arbitrary) hiring threshold s as follows:

$$\lambda(B,s)=\sum_{i\in\{m,w\}} 0.5\int_{B_i}\int_{s}^{\infty} g_i(v\mid a)h_i(a)\Psi_i(a)\,\mathrm{d}v\,\mathrm{d}a. \tag{EC.150}$$

Assume $s_2^*>0$ because otherwise the result trivial follows. Then, from the fixed-point equation (EC.41) (Appendix EC.2.2), s_2^* must be the unique solution to the following equation (with respect to s_2):

$$\lambda(B(s_2,C_2),s_2)=\Delta_2. \tag{EC.151}$$

Building on these observations, we proceed in the following two steps to show that $s_2^*\le\tilde{s}_2$.

Step 1: We first claim that $B_i(s_2,C_2)\subseteq B_i(s_2,\tilde{C}_2)$ for any $\tilde{C}_2>C_2$ and s_2. To see this, because h_i is a strictly positive continuous density function, we observe that $\theta(C_2)$ defined through (EC.148) and (EC.149) strictly decreases in C_2. Hence, we must have $\theta(\tilde{C}_2)<\theta(C_2)$ for any $\tilde{C}_2>C_2$, which implies that $B_i(s_2,C_2)\subseteq B_i(s_2,\tilde{C}_2)$.

Step 2: Using Step 1, we show that $\tilde{s}_2\ge s_2^*$. We first observe that $\lambda(B,s)\le\lambda(\tilde{B},s)$ for any $B\subseteq\tilde{B}$. Hence, from Step 1, we deduce that

$$\lambda(B(s_2^*,\tilde{C}_2),s_2^*)\ge\lambda(B(s_2^*,C_2),s_2^*)=\Delta_2. \tag{EC.152}$$

Further, for any fixed C_2, Lemma EC.1 establishes that $\lambda(B(s_2,C_2),s_2)$ decreases in s_2. Hence, because $\tilde{s}_2$ must satisfy $\lambda(B(\tilde{s}_2,\tilde{C}_2),\tilde{s}_2)=\Delta_2$ (from the fixed-point equation (EC.41)), inequality (EC.152) implies that $\tilde{s}_2\ge s_2^*$. This completes the proof. □

We are now ready to prove (EC.147). Let $\kappa=\sigma_m^2+\tau_m^2=\sigma_w^2+\tau_w^2$ (Assumption 1-(a)). If $C_2=1$, the Firm 2's interview set is simply the entire population. Hence, the match value distribution of interviewees is $v\sim\mathcal{N}(0,\kappa)$, thus $s_2^*=s_2^\rho=\sqrt{\kappa}\Phi^{-1}(1-\Delta_2)$. On the other hand, if $C_2=\Delta_2$, then $s_2^*=s_2^\rho=0$ due to equation (19). By Lemma EC.9, the maximum (respectively, minimum) values of both s_2^* and s_2^ρ are attained when $C_2=1$ (respectively, $C_2=\Delta_2$). Thus, if $s_1^\rho>\sqrt{\kappa}\Phi^{-1}(1-\Delta_2)$, we have $\max_{C_2\in[\Delta_2,1]}\max(s_2^*,s_2^\rho)=\sqrt{\kappa}\Phi^{-1}(1-\Delta_2)<s_1^\rho$ for all $C_2\in[\Delta_2,1]$. Otherwise, we observe that $\min_{C_2\in[\Delta_2,1]}\max(s_2^*,s_2^\rho)=0$. By Lemma EC.8, the end points of the Firm 2's interview set are continuous in C_2, and thus $\max(s_2^*,s_2^\rho)$ is continuous in C_2. Hence, the intermediate value theorem and Lemma EC.9 imply that there exists $\overline{C}_2\ge\Delta_2$ such that $\max(s_2^*,s_2^\rho)\le s_1^\rho$ if $C_2\le\overline{C}_2$. This completes the proof.

EC.5. Beyond Normal Distributions

EC.5.1. Distributions with Increasing Yield Property (Vohra and Yoder 2023)

Vohra and Yoder (2023) consider a similar two-stage hiring model with a single social group, where candidates are characterized solely by their pre-interview scores. Their work focuses on a specific class of conditional match value distributions satisfying the "increasing yield" property (detailed later),

under which the upper interval of Firm 2's interview set is also *greedy* in the score as well. In this section, we show that our meta-characterization recovers the structural results of Vohra and Yoder (2023) as a special case. Specifically, when the increasing yield condition is satisfied, our framework (Proposition EC.1) reproduces the greedy structure of upper interval in Firm 2's interview set.

For simplicity, consistent with Vohra and Yoder (2023), we focus on a single-group setting and thus omit subscript i for groups. We begin by stating the regularity conditions considered in Vohra and Yoder (2023) and the corresponding results.

ASSUMPTION EC.2 **(Regularity Conditions of Vohra and Yoder (2023)).**

(a) ***Hazard rate order:*** *$G(\cdot \mid a)$ is increasing in the hazard order. That is, for $a' > a$ and $s' > s$,*

$$\bar{G}(s' \mid a')\bar{G}(s \mid a) \geq \bar{G}(s \mid a')\bar{G}(s' \mid a),$$

where $\bar{G}(s \mid a) := 1 - G(s \mid a)$.

(b) ***Increasing yield:*** *For all $s_1, s_2 \geq 0$, the function $\left(\int_{s_2}^{\infty} v g(v \mid a)\,\mathrm{d}v\right) \cdot G(s_1 \mid a)$ is increasing in a.*

There are a few conditional match-value distributions that satisfy the increasing yield property. For example, Vohra and Yoder (2023) showed the increasing yield property holds if $v \mid a$ follows an exponential distribution with mean $\lambda(a)$, where $\lambda : \mathbb{R} \to \mathbb{R}$ is an increasing function.

Based on the above assumption, Vohra and Yoder (2023) establish the following result:

COROLLARY EC.2 **(Proposition 3 of Vohra and Yoder (2023)).** *There exists an equilibrium in which:*

1. *Firm 1's optimal interview set is $A_1^* = [a_1^*, \infty)$ for some interview threshold $a_1^* \in \mathbb{R}$*
2. *Firm 2's interview set is $A_2^* = [b_2^*, a_1^*] \cup [c_2^*, \infty)$ for some end points (b_2^*, c_2^*) such that $b_2^* \leq a_1^* \leq c_2^*$.*

We now prove Corollary EC.2 using Proposition EC.1.[27] Because Assumptions EC.1-(i) and (ii) are trivially satisfied, it suffices to verify Assumption EC.1-(iii) by investigating monotonicity of the (discounted) excess value function.

Firm 1: By setting $\Psi(a) = 1$, the optimization problem in (EC.1)-(EC.3) reduces to Firm 1's problem. Let $F(a, s_1) := \mathbb{E}[(v - s_1)_+ \mid a]$ denote the excess value function for a given hiring threshold s_1. Under Assumption EC.2-(a), the function $F(a, s_1)$ is increasing in a for any s_1 (Theorem 1.B.1 of Shaked and Shanthikumar (2007)). This implies that $F(a, s_1)$ satisfies Assumption EC.1-(iii), allowing us to apply Proposition EC.1. By Proposition EC.1, the optimal interview set A_1^* must be a superlevel set of $F(a, s_1^*)$, i.e., $A_1^* = \{a : F(a, s_1^*) \geq \theta\}$ for some threshold θ. Furthermore, since $F(a, s_1^*)$ is strictly increasing in a, this superlevel set must be a single greedy interval of the form $A_1^* = [a_1^*, \infty)$.

Firm 2: Given Firm 1's strategy ($A_1^* = [a_1^*, \infty)$ and s_1^*), the optimization problem in (EC.1)-(EC.3) reduces to Firm 2's problem by setting $\Psi(a) = \mathbb{1}[a < a_1^*] + \mathbb{1}[a \geq a_1^*] G(s_1^* \mid a)$. Let $F_2(a, s_2) :=$

[27] Indeed, we can further apply Proposition EC.2 to conclude that such equilibrium uniquely exists.

$\mathbb{E}[(v-s_2)_+ \mid a]\Psi(a)$ denote the discounted excess value function for a given hiring threshold s_2. Under Assumption EC.2-(a), the function $F_2(a,s_2)$ is increasing in a for $a<a_1^*$ (Theorem 1.B.1 of Shaked and Shanthikumar (2007)). In the following lemma, we establish that $F_2(a,s_2)$ is also increasing for $a>a_1^*$ under Assumption EC.2. Since $F_2(a,s_2)$ satisfies Assumption EC.1-(iii), we can apply Proposition EC.1, which implies that Firm 2's optimal interview set A_2^* must be a superlevel set of $F_2(a,s_2^*)$. Furthermore, since $F_2(a,s_2^*)$ is increasing both for $a<a_1^*$ and for $a>a_1^*$, the superlevel set consists of two disjoint greedy intervals, one below a_1^* and one above it, as desired.

LEMMA EC.10. *Under Assumption EC.2,* $(\int_{s_2}^{\infty}(v-s_2)g(v|a)\,\mathrm{d}v)\cdot G(s_1|a)$ *increases in* a.

Proof of Lemma EC.10. Fix $s_2\geq 0$ and $s_1\geq 0$. We observe Assumption EC.2-(b) holds if and only if

$$\frac{\frac{\partial}{\partial a}\int_{s_2}^{\infty} vg(v|a)\,\mathrm{d}v}{\int_{s_2}^{\infty} vg(v|a)\,\mathrm{d}v}\geq -\frac{\frac{\partial}{\partial a}G(s_1|a)}{G(s_1|a)}. \tag{EC.153}$$

On the other hand, $(\int_{s_2}^{\infty}(v-s_2)g(v|a)\,\mathrm{d}v)\cdot G(s_1|a)$ increases in a if and only if

$$\frac{\frac{\partial}{\partial a}\int_{s_2}^{\infty}(v-s_2)g(v|a)\,\mathrm{d}v}{\int_{s_2}^{\infty}(v-s_2)g(v|a)\,\mathrm{d}v}\geq -\frac{\frac{\partial}{\partial a}G(s_1|a)}{G(s_1|a)}. \tag{EC.154}$$

We claim that

$$\frac{\frac{\partial}{\partial a}\int_{s_2}^{\infty}(v-s_2)g(v|a)\,\mathrm{d}v}{\int_{s_2}^{\infty}(v-s_2)g(v|a)\,\mathrm{d}v}\geq \frac{\frac{\partial}{\partial a}\int_{s_2}^{\infty} vg(v|a)\,\mathrm{d}v}{\int_{s_2}^{\infty} vg(v|a)\,\mathrm{d}v} \tag{EC.155}$$

under Assumption EC.2, which will directly imply our desired result. Indeed, through straightforward algebra, one can show that inequality (EC.155) is true if and only if

$$\frac{\frac{\partial}{\partial a}\int_{s_2}^{\infty} vg(v|a)\,\mathrm{d}v}{\int_{s_2}^{\infty} vg(v|a)\,\mathrm{d}v}\geq \frac{\frac{\partial}{\partial a}\bar{G}(s_2|a)}{\bar{G}(s_2|a)}. \tag{EC.156}$$

From Theorem 1.B.7 of Shaked and Shanthikumar (2007), Assumption EC.2-(a) implies that $\mathbb{E}[v|v\geq s_2,a]$ increases in a. That is,

$$\frac{\partial}{\partial a}\mathbb{E}[v|v\geq s_2,a]=\frac{\partial}{\partial a}\left(\frac{\int_{s_2}^{\infty} vg(v|a)\,\mathrm{d}a}{\bar{G}(s_2|a)}\right) \tag{EC.157}$$

$$=\frac{\left(\frac{\partial}{\partial a}\int_{s_2}^{\infty} vg(v|a)\,\mathrm{d}v\right)\cdot\bar{G}(s_2|a)-\left(\int_{s_2}^{\infty} vg(v|a)\,\mathrm{d}v\right)\cdot\left(\frac{\partial}{\partial a}\bar{G}(s_2|a)\right)}{\bar{G}^2(s_2|a)}\geq 0 \tag{EC.158}$$

which is equivalent to inequality (EC.156). □

EC.5.2. Gumbel Distributions

In Proposition 4 (proved in Appendix EC.4.1), we established that the lower-ranked firm in our model adopts a non-greedy interview strategy, thus forgoing "superstar" candidates (see Appendix EC.4.1 for proofs). As highlighted in Remark 1, this strategy is natural and aligns with anecdotal evidence in labor markets. To further illustrate how such a non-greedy strategy emerges in different settings beyond our

context, we demonstrate that the Gumbel distribution, a foundational model in the consumer choice framework, similarly leads to a non-greedy structure in the strategy of the lower-ranked firm.

We first introduce some notations. We denote the Gumbel distribution with location parameter $a \in \mathbb{R}$ and scale parameter $\beta > 0$ by $\mathsf{Gumbel}(\mu, \beta)$. If a random variable X follows $\mathsf{Gumbel}(\mu, \beta)$, its p.d.f is given by (Ben-Akiva 1985)

$$f(x \mid \mu, \beta) = \exp\left(-\frac{x-\mu}{\beta} - \exp\left(-\frac{x-\mu}{\beta}\right)\right) \tag{EC.159}$$

and its C.D.F. is given by

$$F(x \mid \mu, \beta) = \exp\left(-\exp\left(-\frac{x-\mu}{\beta}\right)\right). \tag{EC.160}$$

It is further known that the mean of $\mathsf{Gumbel}(\mu, \beta)$ is given by $\mathbb{E}[X] = \mu + \beta\gamma$ where $\gamma \approx 0.5772$ is the Euler–Mascheroni constant.

For simplicity, similar to Vohra and Yoder (2023), we consider a single-group setting and omit and subscript i. Recall that we use $g(\cdot \mid a)$ and $G(\cdot \mid a)$ to denote the p.d.f. and C.D.F. of conditional match value $v \mid a$. Suppose that a conditional match value distribution is given by $v \mid a \sim \mathsf{Gumbel}(a - \gamma, 1)$. Here, without loss generality, we standardized the scale parameter as $\beta = 1$ (all of the results naturally extend to $\beta \neq 1$). Note that $\mathbb{E}[v \mid a] = a$, thus the expected match value of a candidate equals to their score (similar to our model with Gaussian distributions).

From Proposition EC.1, characterizing Firm 2's optimal interview set reduces to analyzing the monotonicity of $F_2(a)$.[28] In light of this, we establish in Lemma EC.11 that under the Gumbel distribution, the function $F_2(a)$ exhibits a structure analogous to that illustrated in Figure 3. Specifically, $F_2(a)$ is increasing for $a < a_1^*$ and unimodal for $a > a_1^*$. By Proposition EC.1, this directly implies that Firm 2's optimal interview set follows the same structure as in Proposition 4, namely, a union of a greedy lower interval and a non-greedy upper interval.

Lemma EC.11. *Let $F_2(a) = \mathbb{E}[(v - s_2)_+ \mid a]\Psi(a \mid \boldsymbol{a}_1^*, s_1)$. For any s_1, s_2, and $\boldsymbol{a}_1^*$, the function $F_2(a)$ is increasing for $a < a_1^*$ and unimodal for $a > a_1^*$.*[29]

We prove the above lemma through two claims. First, we show that $F_2(a)$ increases for $a < a_1^*$, where $\Psi(a \mid \boldsymbol{a}_1^*, s_1^*) = 1$. To establish this, it suffices to show that the family of Gumbel distributions is increasing with respect to the location parameter in the sense of the likelihood ratio order (see Fact EC.2 in Appendix EC.3.1).

Claim EC.9. *Let $v \mid a \sim \mathsf{Gumbel}(a - \gamma, 1)$. The family of distributions $\{G(v \mid a)\}_{a \in \mathbb{R}}$ increases in a in the sense of likelihood ratio order. That is, for any $a' > a$, we have $v \mid a' \succeq_{lr} v \mid a$ (see Definition EC.4 for the definition of likelihood ratio order).*

[28] Importantly, the shape of $F_2(a)$ *does not* depend on the score distribution $H(a)$; rather, the score distribution only affects how the threshold level θ is determined in Equation (EC.4).

[29] Since this result holds for *any* s_1 and s_2, it also applies to the optimal hiring thresholds s_1^* and s_2^*.

Proof of Claim EC.9. It suffices to show that, for any $a' > a$, the likelihood ratio $\frac{g(v|a')}{g(v|a)}$ increase in v. We first observe that:

$$\frac{g(v \mid a')}{g(v \mid a)} = \exp\left(a' - a - \exp(a' - v) + \exp(a - v)\right). \tag{EC.161}$$

The above function is increasing in v if and only if a function $q(v) = a' - a - \exp(a' - v) + \exp(a - v)$ is so. Taking derivative of $q(v)$, we have $q'(v) = \exp(a' - v) - \exp(a - v)$ which is positive if and only if $a' > a$. Thus, the likelihood ratio $\frac{g(v|a')}{g(v|a)}$ increases in v, as desired. This completes the proof. □

Next, we show that the function $F_2(a)$ is unimodal for $a > a_1^*$ in the following claim.

Claim EC.10. *Let $v \mid a \sim \mathsf{Gumbel}(a - \gamma, 1)$ and $f(a) := \mathbb{E}[(v - s_2)_+ \mid a] G(s_1 \mid a)$. Then, for any s_1 and s_2, the function $f(a)$ is unimodal.*

Proof of Claim EC.10. We first compute $\mathbb{E}[(v - s_2)_+ \mid a]$ in the following.

$$\mathbb{E}[(v - s_2)_+ \mid a] = \int_{s_2}^{\infty} \bar{G}(v \mid a)\, \mathrm{d}a = \int_{s_2}^{\infty} (1 - \exp(-\exp(a - \gamma - v))\, \mathrm{d}v = \int_0^{\exp(a-\gamma-s_2)} \frac{1 - \exp(-u)}{u}\, \mathrm{d}u. \tag{EC.162}$$

The first equality is by Fact EC.6 (see Appendix EC.3.9). The second equality is from (EC.160). The last equality follows from a variable transformation $u = \exp(a - \gamma - v)$. By differentiating the above equation with respect to a (and after simple algebra), we deduce that

$$\frac{\mathrm{d}\mathbb{E}[(v - s_2)_+ \mid a]}{\mathrm{d}a} = \bar{G}(s_2 | a). \tag{EC.163}$$

Similarly, one can further show that

$$\frac{\mathrm{d}G(s_1 \mid a)}{\mathrm{d}a} = -g(s_1 \mid a). \tag{EC.164}$$

Consequently, the derivative of $f(a)$ is given by

$$\begin{aligned}
f'(a) &= \bar{G}(s_2 \mid a) G(s_1 \mid a) - \mathbb{E}[(v - s_2)_+ \mid a] g(s_1 \mid a) && \text{(EC.165)}\\
&= \bar{G}(s_2 \mid a) g(s_1 \mid a) \left(\frac{G(s_1 \mid a)}{g(s_1 \mid a)} - \mathbb{E}[v - s_2 \mid v - s_2 \geq 0, a] \right) && \text{(EC.166)}\\
&= \bar{G}(s_2 \mid a) g(s_1 \mid a) \underbrace{\left(\exp(-a + \gamma + s_2) - \mathbb{E}[v - s_2 \mid v - s_2 \geq 0, a]\right)}_{:= D(a)} && \text{(EC.167)}
\end{aligned}$$

Because $\lim_{a \to -\infty} D(a) > 0$ and $\lim_{a \to \infty} D(a) < 0$, the intermediate value theorem implies that that there exists a critical point a such that $f'(a) = 0$. We claim that such critical point indeed must be unique, which will imply that the function $f(a)$ is uniomdal. To see why, note that $f(a) = 0$ if and only if $D(a) = 0$, or equivalently

$$\exp(-a + \gamma + s_1) = \mathbb{E}[v - s_2 \mid v - s_2 \geq 0, a]. \tag{EC.168}$$

The left-hand side is clearly a decreasing function in a. For the right-hand side, from Claim EC.9, the random variable $v \mid a$ is increasing in a in the sense of likelihood ratio order. Thus, by Theorem

1.C.6. of Shaked and Shanthikumar (2007), it follows that $\mathbb{E}[v - s_2 \mid v - s_2 \in S, a]$ increases in a for any set S, implying that $\mathbb{E}[v - s_2 \mid v - s_2 \geq 0, a]$ increases in a. Thus, a solution to $D(a) = 0$ must uniquely exist. This completes the proof.

□

EC.6. Competition for Outstanding Candidates in Economics Academic Job Market

In this appendix, we use multi-year data from the EJMR Economics Job Market Wiki (EJMR 2023) to further support our structural results in Proposition 4, in particular the non-greedy structure of the lower-ranked firm. The academic job market in economics serves as a representative example of a setting where schools largely send out interview invitations and hiring offers in parallel, competing for talent from the same pool of candidates. Using data from the economics academic job market, we conduct a simple analysis of how institutions compete for outstanding candidates.

Data. *Economics Job Market Rumors*, also known as EJMR, is an anonymous internet discussion board for economists on the academic job market (EJMR 2023). The website facilitates discussions and information exchange among job market candidates, including an *Economics Job Market Wiki* that tracks the status of the job market.

We focus solely on EJMR's *Economics Job Market Wiki*. We downloaded publicly accessible online data for the academic years 2012-2013 through 2023-2024. Our dataset includes the name of each university, the names and affiliations of candidates invited for on-campus interviews, the list of job offers extended, and the list of accepted job offers.

Data cleaning. For our analysis, we included only faculty openings at U.S. universities, excluding research positions at banks as well as schools and candidates in Europe, Canada, Australia, China, Brazil, and United Arab Emirates. By limiting the dataset to U.S. universities, we aimed to ensure a more reliable comparison across institutions.

To ensure consistency in the dataset, we standardized university and candidate names. University names were often recorded inconsistently, with abbreviations, variations, or misspellings. To address this, we created a university dictionary that converts all recorded university names into their full, standardized form. Similarly, candidate names were sometimes unclear or inconsistent across records. To improve accuracy, we linked identical names and university affiliations, allowing for the precise tracking of candidates throughout the job market process.

Finally, the dataset often contains missing information. We applied logical extrapolation to fill in certain data points where direct reports were partially unavailable. Specifically, if a candidate's name appeared in the accepted offer list, we assumed that they had also received a on-campus interview and a job offer. Likewise, if a candidate received job offers, we inferred that they had also participated in on-campus interviews.

School Tier	U.S. News Rank
Tier 1	1-4
Tier 2	5-10
Tier 3	11-30

Table EC.2 Number schools in each tier based on (U.S. News & World Report 2024)

Finally, we removed data points associated with unknown universities or containing names with numbers or unrecognizable characters were excluded from the dataset.

Data analysis and results. We use the U.S. News & World Report (2024) ranking of economics programs in the U.S. to define five *tiers* of universities. The definition of tiers and the number of schools in each tier are given in Table EC.2. Using this categorization of schools, we define *outstanding candidates* as those who received at least three interviews from Tier 1 affiliation. Given that there are four schools in Tier 1 (Table EC.2), we view these candidates as exceptional in terms of their pre-interview features. In Table EC.3, we show the number of interviews that outstanding candidates received from lower-tier schools. We find that lower-tier schools, in general, tend to interview fewer outstanding candidates. For example, candidate C in the year 2019 (red-colored row in Table EC.3) was interviewed by all of the schools in Tier 1 but by none in Tier 2.

At a high level, this pattern aligns with the structural result established in Proposition 4. Lower-tier schools tend to bypass highly competitive outstanding candidates who secure interviews at top-tier institutions. Our theoretical insight suggests that this behavior arises because the risk of interviewing but ultimately losing these candidates to stronger competitors outweighs the potential benefit of the outstanding candidate's high match value. Notably, this empirical pattern contrasts with Vohra and Yoder (2023), who theorize that firms across all tiers concurrently interview the very top candidates.

Limitations. This data exercise is provided for illustrative purposes and comes with several limitations. First, the wiki is updated through self-reports and other anonymous contributions, leading to a significant number of missing entries. As a result, the data may reflect reporting biases and be skewed (e.g., the set of reported interviews might be incomplete). Second, while the economics job market is largely a representative example of a labor market in which schools make parallel interview and hiring offers, the academic market may not be perfectly vertically differentiated (For example, geographical factors may make a lower tier school more attractive.) Finally, we recognize that EJMR has been heavily criticized for toxic speech in its anonymous online forum (Ederer et al. 2024). We emphasize that our focus is solely on its *Economics Job Market Wiki*, which compiles reports on the job market process and outcomes. We acknowledge these limitations and therefore refrain from drawing definitive conclusions or making recommendations based on this analysis. Rather, we view this exercise as motivation for understanding the rationale behind a non-greedy interview structure of the lower-tier firm in the labor market, as described in Proposition 4.

Year	Candidate ID	# of interviews from Tier 1 (out of 4)	# of interviews from Tier 2 (out of 6)	# of interviews from Tier 3 (out of 20)
2012	A	3	2	0
2012	B	3	2	0
2012	C	3	0	0
2014	A	3	1	1
2015	A	3	2	0
2015	B	3	1	4
2015	C	3	3	2
2015	D	3	1	1
2015	E	3	3	2
2015	F	4	3	3
2016	A	3	2	1
2017	A	3	3	1
2017	B	3	1	0
2017	C	3	5	1
2019	A	3	1	2
2019	B	3	3	1
2019	C	4	0	1
2021	A	3	2	0
2023	A	3	4	1

Table EC.3 Number of Interviews Received by Outstanding Candidates from Each Tier of Schools