FAIR HIRING PROCEDURES

Karl H. Schlag and Andriy Zapechelnyuk

Abstract. In hiring, fair treatment concerns not only the assessment of candidates, but also the process of interviewing and selecting candidates. The latter, so-called procedural fairness, is investigated in this paper using a model of sequential search. We postulate that a hiring procedure is fair if it does not discriminate interviewed candidates by the order in which they are interviewed. We show that procedurally fair hiring prescribes to accept the first candidate who belongs to a prespecified set. Fairness comes at a relatively small cost. The optimal value of hiring without this constraint is at most twice as large.

JEL Classification: D81, D83, J70

Keywords: Sequential search, hiring procedure, fairness, discrimination, equal opportunity, partition strategy, threshold strategy, reservation price, efficiency

Date: July 10, 2023.

Schlag: Department of Economics, University of Vienna, Oskar-Morgenstern-Platz 1, 1090 Vienna, Austria. *E-mail:* karl.schlag@univie.ac.at.
Zapechelnyuk: School of Economics, University of Edinburgh, 31 Buccleuch Place, Edinburgh, EH8 9JT, UK. *E-mail:* azapech@gmail.com.

The authors are grateful for helpful comments to Oliver Fabel, Georgios Gerasimou, Paweł Gola, Toomas Hinnosaar, Maarten Janssen, John Moore, Jakub Steiner, Can Urgun, and participants of various seminars where this paper has been presented.

1. Introduction

Fairness and equal opportunity are of the highest importance in hiring and recruitment. According to Merriam-Webster Dictionary, fair employment constitutes the "employment of workers on a basis of equality without discrimination or segregation especially because of race, color, or creed." Numerous studies investigate the discrimination and statistical biases that may arise when assessing candidates, as surveyed in Bertrand and Duflo (2017) and Neumark (2018). However, there is more to fairness than the concern for how candidates are assessed and compared. The procedure of interviewing and selecting candidates also matters. Any fair assessment of candidates can be undermined if the interviewing process depends not only on the candidates' job related attributes, but also on some of their job unrelated demographic characteristics or the order in which they are interviewed. For instance, this can happen when the person organizing the interviews has a hidden agenda. When not liking a candidate, who would be chosen by the hiring committee, this person might simply decide to interview more candidates until someone more to their liking appears. Principles of fairness need to be formulated to avoid such behavior. In other cases, principles are needed to even understand whether there is unfairness. For instance, a hiring committee might wish to adjust their hiring criteria after conducting several interviews and learning about the quality of the pool of applicants from these interviews. In fact, this form of behavior is recommended by the optimal strategy in the famous secretary problem (Fox and Marnie, 1960). Therein, it is best to interview a certain number of candidates without hiring any of them, and then hire the first candidate who outperforms the initial group. It is our objective to identify whether this is fair.

In this paper we are concerned with unfairness that results from how hiring is organized. We abstract from statistical or behavioral biases by assuming that candidates are assessed fairly. We pursue our investigation within a model of sequential search. Counterfactual scenarios will be used to identify whether hiring is unfair. In particular, we will be looking at who will be hired if candidates would have been interviewed in a different order. We fully characterize hiring rules that are procedurally fair. We then discuss practical implications of procedurally fair hiring, and compute the loss in efficiency due to fairness concerns.

We consider the following model of sequential search.[1] A searcher interviews job candidates one by one in a given order. Each candidate has some publicly observable characteristics and a fit to the job. The job fit is revealed during the interview. It might be measured by the added value. It might incorporate the belonging to a minority group if the searcher takes affirmative action. We are interested in which hiring procedures are fair.

Clearly, procedural fairness is hardly a concern if all candidates are interviewed. However, it is often not feasible or too costly to interview all candidates. Thus, not all candidates can be treated equally as some will be interviewed while others will not. So, we focus our investigation on the fair treatment of those candidates who have been interviewed and whose fit for the job has been revealed. We postulate that interviewed candidates should not be disadvantaged due to the order in which interviews are conducted. Specifically, they should not be hired more likely if they swap their observable characteristics, together with their interview slot, with someone else (possibly, using a multilateral swap). For example, a black person who is interviewed but not hired should not be able to argue that they would have been hired if they had the interview slot and the CV of the white person who got hired.

Our main result is a full characterization of fair hiring procedures. Fair hiring starts with an irrevocable commitment to a categorization of candidates according to their job fit. This has to be done before the interviews. There are three categories: *strong*, *reserve*, and *unacceptable*. The fair hiring principle dictates to hire the first strong candidate on the spot, right after the interview. Reserve candidates can only be hired when everyone has been interviewed and no strong candidate has been found. In that case, any two candidates who have the same job fit must be treated equally. Unacceptable candidates can never be hired.

An implication of procedurally fair hiring is that the hiring standards cannot be manipulated after the interviewing process has started. All criteria that are relevant for the job fit must be discussed and agreed upon in advance. Once the hiring standards are announced and committed to, it is not permitted to include a hidden agenda later. It is also not permitted to "test the waters" by interviewing several

[1]Within our setting, simultaneous search is a special case of sequential search. This is analogous to Burdett and Judd (1983) in the setting of price search.

candidates to make a clearer picture of how good the pool of candidates is, and only then decide how strict the hiring criteria should be.

Another implication is that fair procedures are simple to evaluate and implement. First, the searcher only has to determine for each candidate whether or not they are strong. Second, because the searcher has to commit in advance which fits for the job are categorized as strong, she does not need to update the priors about the pool of candidates, nor to perform optimization calculations after each interview. Third, if the job fit is evaluated by a numerical value, a fair procedure is very simple and is described by a single threshold that separates strong candidates from the others. In this case, finding an optimal fair procedure only involves optimizing over a single parameter, the threshold. This simplicity stands in contrast with the problem of optimization without fairness constraints which is only tractable in special cases.[2]

As the final step, we are left to show how much efficiency is lost if the searcher is required to follow a fair hiring procedure. Here we formally introduce the searcher's optimization problem. The searcher wishes to maximize her expected utility from hiring a candidate when facing the cost of interviews, the discounting of future payoffs, and the uncertainty about candidates. The model of uncertainty is rich in order to add realism. There is systemic uncertainty, modeled by a state of the world that influences the entire pool of candidates. There is also idiosyncratic uncertainty that captures individual differences between the candidates conditional on their observable characteristics. Our only assumption is that candidates who are interviewed earlier are more likely to have higher values than those who are interviewed later, according to first order stochastic dominance. We show that the loss from including the fairness concerns in hiring is not large. Namely, the optimal payoff among all strategies is at most twice the optimal payoff when using fair procedures.

Related Literature. This paper is related to a vast literature on inequality and discrimination in labor markets. The taste-based theory of discrimination goes back to Becker (1957), and the theory of statistical discrimination was founded by Phelps (1972). The up-to-date literature includes numerous theories, empirical studies, and laboratory and field experiments. These studies document and explain pay gaps and other types of inequality and discrimination, propose remedies, and make policy

[2]An optimal search strategy is found under normal distributions with uncertain mean in DeGroot (1968), and under Dirichlet priors in Rothschild (1974).

recommendations. This literature is surveyed in Bertrand and Duflo (2017) and Neumark (2018). Recently, the phenomenon of statistical discrimination received a surge of attention due to the emergence of Big Data and machine learning algorithms. This literature addresses the questions of detection of statistical biases in risk assessment, particularly those emerging from data mining, as well as the design of mechanisms of correction of such biases, using algorithms and machine learning (e.g., Berk, 2012; Berk and Bleich, 2013; Brennan and Oliver, 2013; Feldman et al., 2015; Barocas and Selbst, 2016; Chouldechova, 2017; Corbett-Davies and Goel, 2018). Unlike this entire literature, our paper is not concerned with discrimination due to assessment biases. In fact, we assume that interviews reveal accurate information about the candidates. Instead, we are concerned about procedural fairness.

A major challenge that we had to overcome in this paper is how to define procedural fairness. Labor law regulates hiring procedures, but, as highlighted by Colquitt and Rodell (2015) and Kleinberg et al. (2016), there is no unified view in the law and justice literature on a formal definition of fairness. In the economics literature, a formal treatment of fairness appears in the bargaining and social choice context, where it is captured, among others, by the conditions of envy freeness, anonymity, and symmetry (e.g., Arrow et al., 2010; Vanderschraaf, 2023). This literature is mostly concerned with an equitable allocation of a resource, and the correspondent concept is called distributive or allocative fairness. It is of limited help for us. A hiring procedure cannot be equitable ex-post as everyone wants the job but only one gets it. It would also be unreasonable, inefficient, or even prohibited by law to make it equitable ex-ante, by randomizing between candidates with unequal fit for the job. Note that there is a computer science literature that allows for this kind of randomization to achieve ex-ante fairness (e.g., Kleinberg et al., 2016; Dwork and Ilvento, 2018). The type of fairness that we consider is of intermediate nature. The order of candidates is given, whether it has emerged exogenously or chosen by the searcher. We are highlighting fairness concerns that arise conditional on that order.

We model the process of hiring as sequential search. Sequential search was introduced by Stigler (1962) and McCall (1970), and it has been a tool in an economist's toolkit ever since. Yet the underlying optimization problem, with uncertainty about the pool of candidates and correlation between their values, is complex. Only several special cases have closed form solutions (DeGroot, 1968; Rothschild, 1974; Gastwirth, 1976).

In this paper, when comparing search using fair procedures to general search, we circumvent the complexity of the general search problem by establishing tight bounds for the worst case.

A close relative of sequential search is the secretary problem of Fox and Marnie (1960). The main distinction is that the secretary problem makes no distributional assumptions about the values of candidates. The subsequent literature covers various extensions of this problem as surveyed in Freeman (1983) and Ferguson (1989). More recent papers include adaptations of the secretary problem to online auctions (Babaioff et al., 2008) and to consumer search for products (Parakhonyak and Sobolev, 2015; Schlag and Zapechelnyuk, 2021). Following Babaioff et al. (2008) and Schlag and Zapechelnyuk (2021), we assess the relative cost of fair hiring using the worst-case ratio method, also known as the competitive ratio (Sleator and Tarjan, 1985).

The paper is organized as follows. In Section 2 we introduce the model of hiring as sequential search. In Section 3 we postulate fairness conditions on the hiring procedure, characterize fair procedures, and discuss the insights. In Section 4 we investigate fairness in the special case in which each candidate is given a value. In Section 5 we compute the costs of restricting attention to fair procedures. Section 6 concludes. The proofs are relegated to the online appendix.

2. Model

We use the model of sequential search to formalize a hiring procedure. A searcher interviews job candidates one by one in a given order. Each candidate has some publicly observable characteristics and a fit to the job. The job fit is revealed during the interview. It might be measured by the added value. It might incorporate the belonging to a minority group if the searcher takes affirmative action. We are interested in which hiring procedures are fair.

Specifically, consider a searcher (she) who wishes to hire a candidate (he) from a pool of n job candidates, where $n \geq 2$ ($n = \infty$ is also allowed). The searcher sequentially interviews the candidates, in a given order. The process of conducting interviews and the delay in hiring are costly to the searcher.

Each candidate is identified by his position $i = 1, 2, ..., n$ in the order of interviews, a profile of observable characteristics x_i, and a profile of attributes θ_i. The profile of observable characteristics x_i captures demographics, education, job experience, and other information known about the candidate before the interview. The profile of attributes θ_i captures the fit for the job of candidate i. It only becomes known once the candidate is interviewed. It includes all the information relevant for the searcher's hiring decision, such as productivity, collegiality, and knowledge of the job. We will refer to x_i as candidate i's *CV* and to θ_i as his *type*. Let $\bar{\Theta}$ be a set of possible types.

A profile $X = (x_1, ..., x_n)$ of CVs of all the candidates is given. The order in which candidates are interviewed is also given. There could be different interpretations how this order might have emerged. This could be an exogenous order in which candidates arrive. This could be an order by which the searcher has pre-selected the candidates according to their observable characteristics, possibly based on preference or concerns to improve the outcome. Although the order is given, the possible reasons for this order will later play a role in the interpretation of our fairness criterion.

The type θ_i of candidate i is discovered when this candidate is interviewed. From the searcher's point of view, θ_i is a realization of a random variable conditional on the CV x_i. In addition to this idiosyncratic uncertainty about hidden attributes of the candidates, the searcher has a systemic uncertainty about the pool of candidates as a whole. So the types can be correlated. Our main result does not require any structural assumptions about this uncertainty, so we postpone its formalization until Section 5 where more structure will be needed. Here we only assume that the types of all the candidates are discrete random variables with common support Θ, where $\Theta \subset \bar{\Theta}$. The assumption of discreteness is for clarity of exposition, whereas the assumption of common support is substantive. The implication of the latter is that interviews are valuable. The CVs may provide information about the types of the candidates, but this information is not enough to determine the hidden attributes with certainty. No candidate can be eliminated solely based on his CV.

The searcher has an outside option that can be chosen instead of hiring any candidate. The outside option is available from the beginning, before the first candidate is interviewed, and remains available throughout the hiring process. It is labeled $i = 0$ and identified with a known type θ_0.

The search proceeds in rounds $t = 0, 1, ..., n$. At the outset, in round 0, a type θ_i is drawn for each candidate $i = 1, ..., n$. The searcher (who does not observe the types) decides whether to walk away with the outside option, or to begin the search. In each round $i = 1, ..., n-1$, the searcher decides whether to stop or to continue the search. In round n the search must stop. If the search stops, then the searcher hires one of the interviewed candidates, or chooses the outside option. If the search continues, then the next candidate in the order is interviewed, and so on.

Here we assume that a candidate who is selected to be hired always accepts the job offer. This is for simplicity of exposition. Our main results are not affected if the candidate rejects the offer with a probability that potentially depends on his type.

A hiring strategy prescribes when to stop the search and whom to hire after the search has stopped. It is given by a pair $s = (\sigma, \varphi)$. Let $h_t = (\theta_1, ..., \theta_t)$ be a history of types of interviewed candidates up to round $t = 1, ..., n$, and let h_0 be the empty history. For each round $t = 0, 1, ..., n$, and each possible history h_t in that round, consider the following notation. Let $\sigma(h_t)$ denote the probability of stopping the search in round t. When $t = n$, the search must stop, so set $\sigma(h_n) = 1$. Conditional on stopping in round t, for $i = 1, ..., t$ let $\varphi_i(h_t)$ denote the probability of hiring candidate i, and let $\varphi_0(h_t)$ denote the probability of choosing the outside option, such that $\sum_{i=0}^{t} \varphi_i(h_t) = 1$.

Note that there is free recall, that is, any of the interviewed candidates can be hired. Moreover, our model includes simultaneous search as a special case. Namely, when the searcher decides to interview $k < n$ candidates and then decides which of these to hire, this is as if the searcher interviews these candidates simultaneously.

In what follows, we allow for almost any hiring strategy. In particular, we do not motivate what kind of a strategy the searcher will choose. Hence, we need not introduce any objectives of the searcher at this point of the paper. This allows for a characterization of fairness that is not dependent on the searcher's beliefs or intentions. The only assumption that we make is that there is a candidate who is impossible to beat. Consequently, whenever such a candidate is interviewed, the searcher stops the interviewing process and hires this candidate. Formally, we assume that there exists a type $\bar{\theta} \in \Theta$ such that if a candidate with type $\bar{\theta}$ is interviewed, then the search must stop immediately and this candidate must be hired, so

$$\text{if } \theta_t = \bar{\theta}, \text{ then } \sigma(\theta_1, ...\theta_t) = 1 \text{ and } \varphi_t(\theta_1, ...\theta_t) = 1. \tag{A}$$

We will refer to a candidate with type $\bar{\theta}$ as the *ideal candidate*.

3. Fair Hiring

3.1. **Fairness.** We introduce considerations for procedural fairness to our model of hiring. Observable characteristics may contain information about the likelihood of types of the respective candidates prior to their interviews. This information may influence whether a certain candidate will be interviewed. However, once a candidate is interviewed, then we postulate that this candidate must not be disadvantaged by the observable characteristics and the order of interviews.

To formally define fairness, we compare hiring probabilities when candidates are put in a different order. Thereby, each candidate takes over the CV and the interview slot of the candidate he replaces, and only retains his own type. Changes in orders will be described using permutations or multilateral swaps. Put simply, a swap of two candidates i and j means that candidate i with type θ_i and candidate j with type θ_j swap their CVs and their positions in the order (as if pretending to be each other), while keeping their types. Thus, after the swap, in round i the candidate with CV x_i is interviewed and reveals type θ_j, and in round j the candidate with CV x_j is interviewed and reveals type θ_i. A combination of bilateral swaps is a permutation.

Formally, a permutation function $\pi : \{1, ..., n\} \to \{1, ..., n\}$ is a bijection that maps each candidate $i \in \{1, ..., n\}$ to a candidate $j = \pi(i) \in \{1, ..., n\}$. Let Π be the set of all permutation functions. Given a permutation $\pi \in \Pi$, we denote by $\pi(i)$ the position of candidate i after the permutation is applied. The outside option always remains in the position 0. For notational convenience, let $\pi(0) = 0$.

Given a hiring strategy $s = (\sigma, \varphi)$, a profile of types $r_t = (\hat{\theta}_1, ..., \hat{\theta}_t)$ is called a *run* if, with a strictly positive probability under s, the search reaches round t, the realized types are $(\hat{\theta}_1, ..., \hat{\theta}_t)$, and the search stops in round t, so

$$(\hat{\theta}_1, ..., \hat{\theta}_t) \in \Theta^t \quad \text{and} \quad \left(\prod_{i=0}^{t-1}(1 - \sigma(\hat{\theta}_1, ..., \hat{\theta}_i))\right) \sigma(\hat{\theta}_1, ..., \hat{\theta}_t) > 0.$$

Let $\mathcal{R}_s$ be the set of all runs under strategy s.

Given a permutation $\pi \in \Pi$ and a run $r_t \in \mathcal{R}_s$, let us consider a candidate i who is interviewed in this run, so $i \leq t$. We fix the profile of types $r_t = (\hat{\theta}_1, ..., \hat{\theta}_t)$ revealed in this run, and determine the *ex-ante probability* that candidate i is hired under s conditional on the information about the revealed types. Denote this probability by $q_i(s|r_t)$. We compare $q_i(s|r_t)$ with the probability that the same candidate is hired under s conditional on the same information about candidates' types, but after the candidates are permuted according to π. So, i moves to position $\pi(i)$, and for each $j = 1, ..., t$, the type $\hat{\theta}_j$ is now revealed in round $\pi(j)$. Denote the permuted profile of types by $\pi(r_t)$. The ex-ante probability that candidate in position $\pi(i)$ is hired under s conditional on $\pi(r_t)$ is $q_{\pi(i)}(s|\pi(r_t))$.

Definition 1. A hiring strategy $s = (\sigma, \varphi)$ is *fair* if for every run $r_t \in \mathcal{R}_s$, no permutation of candidates can increase the ex-ante probability that any interviewed candidate is hired or that the outside option is retained:

$$q_i(s|r_t) \geq q_{\pi(i)}(s|\pi(r_t)) \text{ for each } \pi \in \Pi \text{ and each } i \in \{0, 1, ..., t\}. \tag{B}$$

Observe that condition (B) applies not only to each interviewed candidate $i = 1, ..., t$, but also to the outside option $i = 0$. To interpret this, briefly assume that the outside option as a current employee whose contract has come to an end but can be extended. Thus, condition (B) requires that the current employee cannot increase the ex-ante probability of retention by influencing the order in which the candidates are interviewed.

Our theorem below shows that the only hiring strategies that satisfy this fairness criterion are the partition strategies as we now define.

A partition strategy specifies a subset of the possible types of candidates. Candidates whose types are in this subset are called *strong*. Hiring under a partition strategy proceeds as follows. Candidates are sequentially interviewed until either a strong candidate is found, or all candidates have been interviewed. When a strong candidate is found, then interviews stop immediately, and this candidate is hired on the spot. A partition strategy has two additional defining properties when all candidates have been interviewed but none of them were strong. First, the choice of whom to hire must not depend on the order in which the candidates were interviewed. Second, candidates who have the same type must be hired with the same probability.

Definition 2. A hiring strategy $s = (\sigma, \varphi)$ is called a *partition strategy* if there exists a subset of types $\Theta_S \subset \Theta$, with $\bar{\theta} \in \Theta_S$, such that:

(i) For each $t = 1, ..., n-1$ and each $h_t \in \Theta^t$,

$$\sigma(h_t) = \begin{cases} 1 & \text{if } \theta_t \in \Theta_S, \\ 0 & \text{if } \theta_t \notin \Theta_S, \end{cases}$$
$$\varphi_t(h_t) = 1.$$

(ii) For each $h_n \in \Theta^n$,

$$\text{if } \theta_n \in \Theta_S \text{ then } \varphi_n(h_n) = 1,$$
$$\text{for each } i, j \in \{1, ..., n\}, \text{ if } \theta_i = \theta_j \text{ then } \varphi_i(h_n) = \varphi_j(h_n),$$
$$\text{for each } i \in \{1, ..., n\} \text{ and each } \pi \in \Pi, \ \varphi_i(h_n) = \varphi_{\pi(i)}(\pi(h_n)).$$

For interpretation, we further divide the candidates who are not strong into *reserve candidates* and *unacceptable candidates*. Reserve candidates are the candidates who may be hired with positive probability after seeing all candidates, depending on what other types have been revealed during the interviews. Unacceptable candidates are never hired.[3] Thus, when following a partition strategy, the searcher can ask reserve candidates to wait until the search is concluded, whereas she can reject unacceptable candidates on the spot whenever any such candidate is found.

We now present the main result of this paper.

Theorem 1. *A hiring strategy is fair if and only if it is a partition strategy.*

The proof is in Appendix A.1. Here we outline the arguments on an intuitive level.

It is easy to see that every partition strategy is fair. Suppose that an interviewed candidate is strong. Then the search stops and this candidate is hired. Any candidate interviewed earlier is either reserve or unacceptable, and thus can never be hired regardless of the order of interviews. Alternatively, suppose that all candidates have been interviewed and no strong candidate has been found. Then any two reserve candidates with identical types are hired equally likely, and any unacceptable candidate is not hired, independently of the order of interviews. Finally, a partition strategy can choose the outside option only in round 0 (when no one is interviewed),

[3]Formally, a type $\tilde{\theta} \in \Theta$ is reserve if there exists a run $r_t = (\theta_1, ..., \theta_t) \in \mathcal{R}_s$ and $i \in \{1, ..., n\}$ such that $\theta_i = \tilde{\theta}$ and $q_i(s|r_t) > 0$. A type that is neither strong nor reserve is unacceptable.

or in round n (when everyone is interviewed). In either case, no reordering of the candidates can change this outcome. So, condition (B) is satisfied.

Let us sketch the argument for why every fair strategy is a partition strategy. Suppose that the first candidate has been interviewed. For some types of this candidate, the strategy prescribes to stop the search and to hire this candidate with certainty. Let us define such types as strong. Note that the set of strong types is nonempty, because the "ideal candidate" belongs to this set by assumption (A). Several observations are in order. First, if a candidate with a strong type is discovered in any round, the search must stop immediately, and this candidate must be hired. Otherwise, this candidate could have swapped his current position with position 1, where he would have been hired with certainty. This violates condition (B). Second, in every round, the search either stops or continues for sure. Otherwise, if the search continues with an intermediate probability, then a candidate who is sometimes interviewed prefers to be interviewed earlier. This violates condition (B). Third, an order where some reserve candidate can be hired before everyone is interviewed must not exist. To see why, consider a different order in which this reserve candidate is interviewed first, followed by a strong candidate who was previously uninterviewed. In this case the reserve candidate is not hired and would have been better off with the original order. Finally, if the search reaches the last round, candidates with equal types must be treated equally, as otherwise one would want to swap the position with the other. Moreover, who is hired must not depend on the order of interviews, as otherwise condition (B) is violated. Summarizing the above, we conclude that a hiring strategy that satisfies property (B) must be a partition strategy.

3.2. **Discussion.** Following Theorem 1, fair hiring means using partition strategies. In this section we outline some implications of using such strategies and discuss some assumptions.

We first highlight several implications.

Commitment to the hiring criterion. Before starting the interviewing process, the hiring committee has to agree on a categorization of the candidates' attributes into who will be hired on the spot, who will be put in reserve, and who will never be hired. The hiring committee also has to agree that this categorization may not be changed, regardless of what happens during the interviews. In particular, fair hiring precludes the use of random events (like flipping a coin) during the interviews. Randomness may only be used as a tie-breaker between candidates of equal types once all candidates have been interviewed.

Flexibility of objectives. When implementing a fair hiring procedure, the searcher must categorize the candidates' types, but she is free to choose which types belong to which category. The categorization can depend the searcher's objective function, her costs, and her outside option. For example, the searcher may wish to hire quickly, in which case she would label many types as strong. At the other extreme, the searcher may wish to interview all candidates unless an unusually excellent candidate is found. In this case she can categorize only very few candidates as strong.

Affirmative action. Affirmative action may not be implemented by choosing a minority group candidate over another candidate when both have the same type. It has to be implemented by incorporating the attribute of belonging to the minority group into the type. In that case, candidates can be assigned to different categories based on whether they belong to the minority group.

Testable implications. An outsider might not be able to observe the hiring procedure that the searcher follows. Yet, one can still identify that the hiring procedure is unfair when the search stops before all the candidates are interviewed, and the last interviewed candidate is not hired.

Individual treatment. Candidates have to be interviewed one by one with a decision being made after each interview. In particular, the searcher may not first wait to see the first few candidates before making a decision.

No learning. Learning about the quality of the pool of candidates is not allowed. In particular, Bayes' rule is not used when following a partition strategy. Even if many "bad" candidates are interviewed, the searcher may not get disappointed and stop the search to hire a candidate who was initially categorized as reserve. Similarly, when seeing a "good" candidate, the searcher may not get excited about the pool and decide to look further if that candidate was initially categorized as strong.

Next, we discuss several assumptions made in our model.

Fixed order of interviews. One might wish to mitigate the unfairness that arises due to the order in which candidates are interviewed by randomizing the order. However, randomization can easily increase a feeling of injustice and is hardly realistic as this creates inefficiencies. In practice, candidates are often interviewed in the order determined by the hiring committee based on the candidates' CVs, or in the order in which applications arrive. So, the order is often predetermined. Even if the order is random, the fairness concerns raised in this paper are still relevant as they apply conditional on each possible realization of the order.

Concern for interviewed candidates. Fairness concerns raised in this paper are directed towards possible complaints of interviewed candidates. These complaints arise from objecting that a hypothetical alternative order would have increased their hiring probability. Uninterviewed candidates are not allowed to complain. This gives rise to the question: What if we apply the same fairness concern to all candidates, interviewed and uninterviewed? Then this will mandate to interview all candidates in the pool. While this is clearly fair, it might be too restrictive in practice if the cost of interviewing is high.

Common support of types. Ex-ante the searcher is endowed with beliefs about the candidates' types conditional on their CVs. The assumption that these beliefs have common support reflects an implicit concern for fairness. This means that the searcher cannot be certain that one candidate is better than another just by looking at their CVs. It is always possible that the revealed performance in the interview outweighs the negative expectations that emerge from a bad CV.

Perfect recall. We have assumed that the searcher has perfect recall, namely, that any candidate previously interviewed is still available for hiring when the search stops. Yet under a partition strategy the searcher will not need this recall option unless all candidates have been interviewed.

No rejections. For simplicity of exposition, we have assumed that a selected candidate always accepts the job offer. One can easily incorporate the possibility that a job offer can be rejected. In that case, the searcher treats this candidate as if his type was unacceptable, and continues to follow the partition strategy. In particular, fair hiring means that when a selected candidate turns down an offer, then the search continues. The offer may not be passed on to a candidate interviewed in the past, unless all candidates have been interviewed.

4. Fair Hiring Using Thresholds

In this section we make some assumptions on the objective of the searcher. This restricts the kinds of search strategies she would use. We investigate what this implies for fair hiring.

Assume that the searcher's preferences over candidates' types admit a utility representation. Let $v(\theta)$ be the searcher's added value from hiring a candidate with type θ as compared with the outside option. Thus, the value of the outside option is zero, so $v(\theta_0) = 0$. We remain agnostic as to how the searcher evaluates future outcomes and incorporates costs of interviewing.

We are interested in fair hiring when the searcher evaluates the candidates' types according to $v(\cdot)$. Following Theorem 1, a fair hiring strategy has to be a partition strategy. The first insight is that any candidate with a negative value is unacceptable. Observe now what can happen with a candidate whose value is positive using the following example. Let the set of types be $\Theta = \{\text{mediocre}, \text{good}, \text{excellent}\}$ and let

$$v(\text{mediocre}) = 1,\ v(\text{good}) = 2,\ v(\text{excellent}) = 3.$$

Consider a partition strategy in which mediocre and excellent candidates are categorized as strong, and good candidates are categorized as reserve. This is a fair hiring strategy according to Theorem 1. However, according to this strategy, after first interviewing a good candidate followed by a mediocre candidate, the interview process stops and the mediocre candidate is hired. Such behavior is not optimal, as the searcher strictly prefers the good candidate to the mediocre one. Furthermore, it does not seem fair that the mediocre candidate is hired when the good candidate has been interviewed. To rule out such situations, we introduce a new condition. This condition can be interpreted both as a constraint that follows from optimality and as a constraint that is imposed to due to fairness considerations.

In what follows, we assume that the search strategy satisfies the following condition. Only the candidates with the highest value can be hired, where ties are broken equally.

Formally:

$$\begin{aligned}&\text{For each run } r_t = (\theta_1, ..., \theta_t) \in \mathcal{R}_s \text{ and each } i, j \in \{0, 1, ..., t\},\\&\text{if } v(\theta_i) < v(\theta_j) \text{ then } \varphi_i(r_t) = 0, \text{ and}\\&\text{if } v(\theta_i) = v(\theta_j) \text{ then } \varphi_i(r_t) = \varphi_j(r_t).\end{aligned} \tag{C}$$

We now show that strategies that are both fair and satisfy (C) are the threshold strategies as we now define.

A threshold strategy is a particular partition strategy whose partition can be described by a threshold. This threshold is weakly above the value of the outside option. The candidates whose values are at least as high as the threshold are considered strong.[4] The rest of the candidates are divided into reserve and unacceptable. Reserve candidates have values below the threshold but above the outside option. Unacceptable candidates have values below the outside option. Candidates who are as good as the outside option can be considered either reserve or unacceptable. Hiring proceeds as follows. Candidates are sequentially interviewed until either a strong candidate is found, or all candidates have been interviewed. When a strong candidate is found, then interviews stop immediately, and this candidate is hired on the spot.[5] If instead all candidates are interviewed and no strong candidate is found, then the following happens. The best reserve candidate is hired provided he is better than the outside option. If the best reserve candidate is as good as the outside option, then possibly no one is hired. If there are several best candidates, then they are hired with equal probabilities. If all the candidates are unacceptable, then no one is hired.

Definition 3. A partition strategy $s = (\sigma, \varphi)$ is called a *threshold strategy* if there exists a threshold $y \in \mathbb{R}_+$ and a partition $\{\Theta_S, \Theta_R, \Theta_U\}$ of Θ such that for each $\theta \in \Theta$,

$$\theta \in \begin{cases} \Theta_S & \text{if } v(\theta) \geq y, \\ \Theta_R & \text{if } 0 < v(\theta) < y, \\ \Theta_R \cup \Theta_U & \text{if } v(\theta) = 0, \\ \Theta_U & \text{if } v(\theta) < 0, \end{cases}$$

[4]Without loss of generality, the candidates whose value is equal to the threshold are considered strong. This is because we have assumed that the types have a discrete support.

[5]Strategies based on this principle are also known in the search literature as reservation price strategies.

and for each $h_n = (\theta_1, ..., \theta_n) \in \Theta^n$,

$$\sum\nolimits_{i \in \arg\max_{j=0,1,...,n} v(\theta_j)} \varphi_i(h_n) = 1,$$

$$\text{for each } i, j \in \{1, ..., n\}, \text{ if } v(\theta_i) = v(\theta_j) \text{ then } \varphi_i(h_n) = \varphi_j(h_n).$$

Corollary 1. *A fair strategy satisfies condition* (C) *if and only if it is a threshold strategy.*

We show why Corollary 1 follows easily from Theorem 1, namely, that partition strategies that satisfy condition (C) are threshold strategies. Suppose that the searcher interviews a reserve candidate A followed by a strong candidate B. The partition strategy dictates that the searcher must stop and hire B with probability one. Condition (C) requires that B can only be hired with probability one if B's value is greater than A's. This leads to the conclusion that the value of every strong candidate must be greater than that of every reserve candidate. Condition (C) also implies that when all candidates are interviewed, only the best of them can be hired.

We hasten to point out the following implication.

Remark 1. When the searcher is restricted to threshold strategies, her optimization problem is greatly simplified. There is a single variable that the searcher needs to choose, namely, the threshold.

5. Cost of Fairness

Now let us address the question of how much potential value the searcher loses by restricting herself to fair strategies. To do this, we introduce the searcher's utility that incorporates time preferences and interviewing costs. We also specify her beliefs about the candidates' types. We compare the maximum expected utility of the searcher with and without the constraint to use fair strategies. We perform this investigation in a special but relevant case where earlier candidates are more likely to have higher values.

First we introduce the searcher's utility. Following Section 4, each type θ of a candidate is associated with a value $v(\theta)$. We make the following additional assumptions. The searcher incurs a cost $c \geq 0$ for each interview. Future payoffs are converted into their present value using a discount factor $\delta \in (0,1)$. Thus, if the searcher stops the search in round t and hires a candidate with a type θ, her utility is specified to be

$$u(\theta,t) = \begin{cases} u_0 & \text{if } t = 0, \\ u_0 + \delta^t v(\theta) - \sum_{k=1}^{t} \delta^k c & \text{if } t \geq 1, \end{cases}$$

where u_0 is the baseline utility that the searcher obtains if she decides to stop in round $t = 0$ and choose the outside option.

We assume that the baseline utility u_0 is enough to cover the cost of interviewing the entire pool of applicants, so

$$\sum_{t=1}^{n} \delta^t c \leq u_0. \tag{1}$$

Thus, the net utility of the searcher cannot go below zero. A possible interpretation is that the searcher is a firm owner, and u_0 is an initial value of the firm. Assumption (1) means that the firm does not need to borrow in order to pay the costs of its hiring process. For a given baseline utility $u_0 > 0$, let C_{u_0} be the set of pairs of the interview cost c and discount factor δ that satisfy the above constraint, so

$$C_{u_0} = \left\{ (c,\delta) \in \mathbb{R}_+ \times (0,1) : \sum\nolimits_{t=1}^{n} \delta^t c \leq u_0 \right\}.$$

The searcher's uncertainty about the candidates' types is modeled as follows. Let Ω be a finite set that describes the possible states of the world. The searcher has a prior β over Ω. In each state $\omega \in \Omega$, the type θ_i of candidate $i = 1, ..., n$ is distributed according to a discrete probability distribution $\lambda_\omega(\cdot|x_i)$ that is conditional on his observable CV x_i. Thus, in each state ω, the types are distributed independently but not necessarily identically. For all $x_i \in (x_1, ..., x_n)$ and all $\omega \in \Omega$, the distributions $\lambda_\omega(\cdot|x_i)$ are assumed to have common support $\Theta \subset \bar{\Theta}$. Prior β is described by a profile that assigns a probability p_ω to each distribution λ_ω, so $\beta = (\lambda_\omega, p_\omega)_{\omega \in \Omega}$.

Let $F_\omega(\cdot|x_i)$ be the cumulative probability distribution of values conditional on x_i in state ω, so

$$F_\omega(\hat{v}|x_i) = \Pr\left(v(\theta) \leq \hat{v} \middle| \omega, x_i\right) = \int_{\theta \in \Theta : v(\theta) \leq \hat{v}} \lambda_\omega(\mathrm{d}\theta|x_i). \tag{2}$$

We assume that candidates who are positioned earlier in the order are likely to have higher values in the following sense. For each state ω, the conditional distributions of the values of earlier candidates first order stochastically dominate those of later candidates, so

$$\text{for all } \omega \in \Omega,\ F_\omega(\cdot|x_i) \succeq_{fosd} F_\omega(\cdot|x_j) \text{ whenever } i < j. \tag{3}$$

Let $\mathcal{B}_{fosd}$ be the set of priors that satisfy condition (3).

The story behind condition (3) is as follows. Assume briefly that there is no given order of candidates, but that the candidates are ranked according to their CVs as follows. Independently of the state of the world, the value distribution of a candidate with a better CV first order stochastically dominates that of a candidate with a worse CV. Then, following Weitzman (1979), it is best to interview the candidates with higher ranked CVs first. This leads to an order that satisfies condition (3).

We now define and compare optimal payoffs with and without constraints imposed by fair hiring. Let $U_s(\beta, c, \delta)$ denote the expected utility of the hiring strategy s given prior $\beta \in \mathcal{B}_{fosd}$, cost parameter c, and discount factor δ. Let $U^*(\beta, c, \delta)$ be the maximal expected payoff when choosing among all hiring strategies. Let $U^P(\beta, c, \delta)$ and $U^T(\beta, c, \delta)$ be the maximal expected payoffs when choosing among partition and threshold strategies, respectively.

We are now able to measure how much the searcher loses by limiting herself to fair hiring strategies. As a measure, in accordance with Theorem 1, we look at the ratio $U^*(\beta, c, \delta)/U^P(\beta, c, \delta)$. This ratio describes how many times the searcher could have improved her payoff if she is not constrained to use fair strategies. We will refer to this ratio as the *relative cost of fairness*. We also do as above for the case where the hiring strategy has to additionally satisfy condition (C). In accordance with Corollary 1, we evaluate $U^*(\beta, c, \delta)/U^T(\beta, c, \delta)$. This ratio is referred to as the *relative costs of fairness under* (C).

An obstacle that stands in our way is that it is generally intractable to compute $U^*(\beta, c, \delta)$. This value can only be found under very specific assumptions about the prior β.[6] This means that it is also generally intractable to compute the relative cost of fairness for general priors. Yet we are able to provide tight upper bounds on the relative cost of fairness.

The next theorem provides tight upper bounds on the relative cost of fairness, with and without condition (C). It turns out that both of these bounds are equal to 2.

[6]These include i.i.d. normal distribution with uncertain mean (DeGroot, 1968) and a Dirichlet prior (Rothschild, 1974).

Theorem 2.

$$\sup_{\beta \in \mathcal{B}_{fosd},\,(c,\delta)\in C_{u_0}} \frac{U^*(\beta,c,\delta)}{U^P(\beta,c,\delta)} = \sup_{\beta \in \mathcal{B}_{fosd},\,(c,\delta)\in C_{u_0}} \frac{U^*(\beta,c,\delta)}{U^T(\beta,c,\delta)} = 2.$$

The proof is in Appendix A.2. Here we outline the intuition behind the proof.

The first key step to the proof is to find an upper bound on the relative cost of fairness whose supremum can be evaluated. It is important that this upper bound is tight, so that in the worst case the upper bound is approximately equal to the original expression. The upper bound is found by replacing the unconstrained searcher (whose payoff $U^*(\beta,c,\delta)$ is in the numerator) by an unconstrained partially informed searcher who observes the state ω. Because the conditional distribution $F_\omega(\cdot|x_t)$ is decreasing in t by first order stochastic dominance, the optimal expected payoff of this strategy can be evaluated using Weitzman (1979).

The next steps consist of showing that the ratio $U^*(\beta,c,\delta)/U^T(\beta,c,\delta)$ is the largest when there are only two types. This is then used to explicitly bound this ratio by 2. An example shows that this bound is tight. This completes the proof for the relative cost of fairness under (C).

In the final step, this is connected to the cost of fairness without (C). While the maximal expected payoff among the partition strategies is generally greater than that among the threshold strategies, we show that the two are equal when considering only distributions involving two types. This shows that the two bounds are equal and tight.

The intuition for why $U^*(\beta,c,\delta)/U^T(\beta,c,\delta) \leq 2$ is as follows. When the searcher uses a threshold strategy, she lacks flexibility to adapt the threshold in different states of the world and in different rounds of the search. In some states or rounds, the searcher would prefer to stop when her threshold strategy prescribes to keep searching. In other states or rounds, the searcher would prefer to keep searching when her threshold strategy prescribes to stop. As there are only two ways to get it wrong, and the payoffs cannot be negative by assumption (1), balancing the two concerns yields the error factor of at most 2.[7]

[7]Despite some similarities, our result is not a variation of the prophet inequality of Krengel and Sucheston (1978). In the prophet inequality setting, the prophet knows which candidate is the best. The bound of 2 is then obtained for the case of independent types. In contrast, in our setting, types can be correlated, and the partially informed searcher only knows the state ω. Given our assumption (2) of the first order stochastic dominance, knowing the state is not enough to determine which candidate is the best.

Observe that none of the parameters, c, δ, or u_0, affect the upper bound on the relative cost of fairness. This is because the case where the restriction to threshold strategies hurts the most is when the potential values of the candidates are so high that the value of outside option u_0 is negligible. In this case, by assumption (1), the maximum accumulated cost of search, $\sum_{t=1}^{n} \delta^t c$, is also negligible. As the costs do not matter, and both $U^*(\beta, c, \delta)$ and $U^T(\beta, c, \delta)$ are affected by the discount factor δ in the same way, δ does not matter either.

6. Concluding Remarks

There is more to fairness than only the concern for how candidates are evaluated and compared. The procedure for selecting among the candidates also matters. Any fair evaluation of candidates can be undermined by strategically manipulating the interviewing procedure. We model the hiring process as sequential search, and use this to postulate fairness criteria on how this process should be conducted. Interestingly, only very simple hiring strategies turn out to be fair. In the canonical search model where candidates are given values, these fair strategies are the threshold strategies. Threshold strategies are found to be optimal in the particular setting where the values are i.i.d. (e.g., McCall, 1970). Our investigation shows that hiring strategies must be restricted to that class if fairness is to be taken into account. This means that threshold strategies are also relevant under more general distributions. In particular, the restriction to threshold strategies substantially simplifies finding best hiring strategies under more realistic distributions.

Appendix

A.1. **Proof of Theorem 1.** Consider any strategy s satisfies (A). By our argument after Theorem 1, it is straightforward to verify that if s is a partition strategy, then it satisfies (B). We now prove that if s satisfies (B), then s is a partition strategy.

Let $s = (\sigma, \varphi)$ be a strategy that satisfies (A) and (B). We show that s is the partition strategy whose set of *strong* types, denoted by Θ_S, is given as follows. Let Θ_S be the subset of types in Θ such that the search stops in round 1 and the first interviewed candidate is hired with certainty, so

$$\Theta_S = \{\theta_1 \in \Theta : \sigma(\theta_1) = 1 \text{ and } \varphi_1(\theta_1) = 1\}. \tag{4}$$

Note that Θ_S is nonempty, because the ideal type $\bar{\theta}$ is in Θ_S by condition (A).

Let σ_0 be the probability that the search stops in round 0, so $\sigma_0 = \sigma(h_0)$. Throughout the proof, we will be assuming that

$$\sigma_0 < 1. \tag{5}$$

Otherwise, the search stops with certainty in round 0, so no one is interviewed, which trivially satisfies the definition of the partition strategy. The rest of the proof is divided into four steps.

Step 1. *Whenever a strong candidate is interviewed, this candidate must have probability one to be hired. Moreover, the search must have continued with certainty after interviewing earlier candidates. Formally, for each run $r_t = (\theta_1, ..., \theta_t) \in \mathcal{R}_s$ and each $i \in \{1, ..., t\}$, if $\theta_i \in \Theta_S$, then*

$$\sigma(\theta_1, ..., \theta_i) = 1,\ \varphi_i(\theta_1, ..., \theta_i) = 1, \text{ and } \sigma(\theta_1, ..., \theta_j) = 0 \text{ for each } j = 1, ..., i-1. \tag{6}$$

Proof of Step 1. By contradiction, suppose that there exist a run $r_t \in \mathcal{R}_s$ and a candidate in position $i^B \in \{1, ..., t\}$ with a type $\theta_{i^B} \in \Theta_S$ such that (6) does not hold. Denote this candidate by B and her type by θ^B, so $\theta^B = \theta_{i^B}$. The ex-ante probability that B is hired conditional on run r_t is given by

$$q_{i^B}(s|r_t) = (1-\sigma_0)\left(\prod_{j=1}^{i^B-1}(1-\sigma(\theta_1, ..., \theta_j))\right)\sigma(\theta_1, ..., \theta_{i^B})\varphi_{i^B}(\theta_1, ..., \theta_{i^B}) < 1-\sigma_0,$$

where the inequality is by (5), and because we have assumed that (6) does not hold.

Consider the permutation π that swaps i^B and 1. Thus, candidate B is interviewed in position $\pi(i^B) = 1$ after the permutation. By $\theta^B \in \Theta_S$ and (4), candidate B is hired with certainty in round 1, so the ex-ante probability that B is hired is given by

$$q_{\pi(i^B)}(s|\pi(r_t)) = (1-\sigma_0)\sigma(\theta^B)\varphi_{\pi(i^B)}(\theta^B) = 1-\sigma_0.$$

We thus obtain $q_{i^B}(s|r_t) < q_{\pi(i^B)}(s|\pi(r_t))$, which contradicts condition (B). This completes the proof of Step 1. □

Step 2. *In each round, the search either stops or continues with probability 1. Formally, for each run $r_t = (\theta_1, ..., \theta_t) \in \mathcal{R}_s$, and each $i \in \{1, ..., t\}$,*

$$\sigma(\theta_1, ..., \theta_i) \in \{0, 1\}.$$

Proof of Step 2. By contradiction, suppose that there exist a run $r_t = (\theta_1, ..., \theta_t) \in \mathcal{R}_s$ and a round $i \in \{1, ..., t\}$ such that $\sigma(\theta_1, ..., \theta_i) \in (0, 1)$. Because the search stops in round i with a positive probability, $r_i = (\theta_1, ..., \theta_i)$ is also a run in R_s.

Let $\hat{\theta}_{i+1} \in \Theta_S$. Observe that $\hat{r}_{i+1} = (\theta_1, ..., \theta_i, \hat{\theta}_{i+1})$ is a run in $\mathcal{R}_s$, because r_i is a run, the search continues in round i with a positive probability, and by Step 1 it stops in round $i+1$ under $\hat{r}_{i+1}$. Then, by Step 1 and the assumption that $\hat{\theta}_{i+1} \in \Theta_S$, we conclude that $\sigma(\theta_1, ..., \theta_i) = 0$. Thus, we have reached a contradiction. □

Step 3. *The search cannot stop after a candidate who is not strong is interviewed, unless this is the last candidate. Formally, for each* $t = 1, ..., n-1$*, each run* $r_t = (\theta_1, ..., \theta_t) \in \mathcal{R}_s$*, and each* $i \in \{1, ..., t\}$*,*

$$\text{if } \theta_i \notin \Theta_S, \text{ then } \sigma(\theta_1, ..., \theta_i) = 0.$$

Proof of Step 3. By contradiction, suppose that there exist a round $t < n$, a run $r_t = (\theta_1, ..., \theta_t) \in \mathcal{R}_s$, and a candidate $i \in \{1, ..., t\}$ with a type $\theta_i \notin \Theta_S$ such that $\sigma(\theta_1, ..., \theta_i) > 0$. By Step 2 it means that

$$\sigma(\theta_1, ..., \theta_i) = 1, \quad \text{and } \sigma(\theta_1, ..., \theta_j) = 0 \text{ for each } j = 1, ..., i-1.$$

In particular, the search stops in round i. That is, i is the last round of the run r_t, so $t = i$ and $r_t = r_i = (\theta_1, ..., \theta_i)$.

First, suppose that once the search stops in round i, the outside option is chosen with certainty, so $\varphi_0(\theta_1, ..., \theta_i) = 1$. Observe that

$$q_0(s|r_i) = 1,$$

because the search does not stop in rounds $1, ..., i-1$, and the outside option is chosen in either round 0 or round i. Denote the candidate in position $i+1$ by C and her type by θ^C, so $\theta^C = \theta_{i+1}$. Consider the permutation π that assigns C to position 1, and assigns the rest of the candidates arbitrarily. Denote by π^{-1} the inverse permutation.

Let $\theta^C \in \Theta_S$. Consider the sequence with a single interview given by $\hat{r}_1 = (\theta^C)$. Because $\theta^C \in \Theta_S$, by Step 1 we have $\sigma(\theta^C) = 1$. Thus, $\hat{r}_1$ is a run in $\mathcal{R}_s$. Moreover, by Step 1, candidate C is hired with certainty, provided that C is interviewed. Thus, the outside option cannot be chosen with probability one under $\hat{r}_1$, so $q_0(s|\hat{r}_1) = \sigma_0$, and we have σ_0. Now, apply the permutation π^{-1} conditional on the run $\hat{r}_1$. Under this permutation, we restore the original order where C is in the position $i+1$. Under this order, we have $q_0(s|r_i) = 1$ as we established earlier. In summary, we obtain $q_0(s|\hat{r}_1) = \sigma_0 < 1 = q_0(s|r_i)$, where the inequality is by (5). Thus we have reached a contradiction with condition (B).

Second, suppose that once the search stops in round i, the outside option is not chosen with certainty, so $\varphi_0(\theta_1, ..., \theta_i) < 1$. Let $i^B \in \{1, ..., i\}$ be a candidate who may be hired in round i, so $\varphi_{i^B}(\theta_1, ..., \theta_i) > 0$. Denote this candidate by B and her type by θ^B, so $\theta^B = \theta_{i^B}$. Denote the candidate in position $i+1$ by C and her type by θ^C, so $\theta^C = \theta_{i+1}$. Consider the permutation π that assigns B to position 1 and C to position 2, and assigns the rest of the candidates arbitrarily. Denote by π^{-1} the inverse permutation.

Let $\theta^C \in \Theta_S$, and consider the sequence $\hat{r}_2 = (\theta^B, \theta^C)$. Because $\theta^B \notin \Theta_S$, by (4) we have $\sigma(\theta^B) > 0$. Because $\theta^C \in \Theta_S$, by Step 1 we have $\sigma(\theta^B, \theta^C) = 1$. Thus, $\hat{r}_2$ is a run in $\mathcal{R}_s$. Moreover, by Step 1, candidate C is hired with certainty. Thus, candidate B (who is in position 1) cannot be hired, so $q_1(s|\hat{r}_2) = 0$.

Now, apply the permutation π^{-1} conditional on the run $\hat{r}_2$. Under this permutation, we restore the original order where B and C are in positions i^B and $i+1$, respectively. Under this order, there is a strictly positive probability that run r_i occurs, in which case B may be hired, so $q_{i^B}(s|\pi^{-1}(\hat{r}_2)) > 0$. We thus obtain

$$q_{i^B}(s|\pi^{-1}(\hat{r}_2)) > q_1(s|\hat{r}_2) = 0,$$

which contradicts condition (B). This completes the proof of Step 3. □

Step 4. *If all the candidates are interviewed, then hiring probabilities do not depend on the order, and candidates with the same type are hired equally likely. Formally, for each run $r_n = (\theta_1, ..., \theta_n) \in \mathcal{R}_s$ and each $i \in \{1, ..., n\}$,*

$$\text{for each } j \in \{1, ..., n\}, \text{ if } \theta_i = \theta_j \text{ then } \varphi_i(\theta_1, ..., \theta_n) = \varphi_j(\theta_1, ..., \theta_n), \tag{7}$$

and

$$\text{for each } \pi \in \Pi,\ \varphi_i(r_n) = \varphi_{\pi(i)}(\pi(r_n). \tag{8}$$

Proof of Step 4. Let $r_n = (\theta_1, ..., \theta_n)$ be a run in $\mathcal{R}_s$. Condition (8) is straightforward by (B). We now show (7).

Suppose first that $\theta_n \in \Theta_S$. By Steps 1 and 2, for all $t = 1, ..., n-1$ we have $\theta_t \notin \Theta_S$, and thus $\theta_t \neq \theta_n$. By Step 1, $\varphi_n(\theta_1, ..., \theta_n) = 1$. Therefore, $\varphi_t(\theta_1, ..., \theta_n) = 0$ for all $t = 1, ..., n-1$, and hence (7) holds.

Alternatively, suppose that $\theta_n \notin \Theta_S$. By Steps 1 and 2, for each $t = 1, ..., n-1$ we have $\theta_t \notin \Theta_S$ and $\sigma(\theta_1, ..., \theta_t) = 0$, so round n is reached with certainty conditional on the run r_n. Consequently, for each $t = 1, ..., n$ the ex-ante probability of being hired conditional on r_n is given by

$$q_t(s|r_n) = (1-\sigma_0)\varphi_t(\theta_1, ..., \theta_n). \tag{9}$$

By (5) we have $\sigma_0 < 1$, as otherwise r_n would not be a run in $\mathcal{R}_s$.

By contradiction, suppose that there exist $i, j \in \{1, ..., n\}$ such that $\theta_i = \theta_j$ but $\varphi_i(\theta_1, ..., \theta_n) < \varphi_j(\theta_1, ..., \theta_n)$. Using (5) and (9), this implies that

$$q_i(s|r_n) < q_j(s|r_n). \tag{10}$$

Now, consider the permutation π that swaps i and j. Because $\theta_i = \theta_j$, the run $r_t = (\theta_1, ..., \theta_n)$ is unchanged under this permutation, so

$$q_t(s|\pi(r_n)) = q_t(s|r_n) \quad \text{for each } t = 1, ..., n. \tag{11}$$

Thus we obtain

$$q_{\pi(i)}(s|\pi(r_n)) = q_j(s|\pi(r_n)) = q_j(s|r_n) > q_i(s|r_n),$$

where the first equality is by $\pi(i) = j$, the second equality is by (11), and the inequality is by (10). We have reached a contradiction with condition (B). This completes the proof of Step 4. □

Steps 1–4 imply that every strategy that satisfies (A) and (B) is a partition strategy. Indeed, by the definition of the set Θ_S of strong types and Step 1, the procedure stops and hires the last interviewed candidate when that candidate is strong. By Step 3, whenever a candidate who is not strong is interviewed, the searcher does not stop unless all candidates have been interviewed. By Step 4, when all candidates have been interviewed, the candidates with equal types are treated equally. This completes the proof of Theorem 1.

A.2. **Proof of Theorem 2.** We will use the following notation throughout the proof. Given a prior $\beta = (\lambda_\omega, p_\omega)_{\omega\in\Omega}$, let Θ_β be the common support of the distributions over types under prior β. Recall that the distributions are discrete, so Θ_β is finite or countable. For each ω, let F_ω be the conditional distribution of values induced by the distribution of types λ_ω according to (2).

First, we introduce the notion of strict threshold strategy. A strict threshold strategy is similar to a threshold strategy except in the event that all the candidates have been interviewed and all of them are below the threshold, this strategy selects the outside option.

Definition 4. A threshold strategy $s = (\sigma, \varphi)$ is *strict* if for each $(\theta_1, ..., \theta_n) \in \bar{\Theta}^n$,

$$\text{if } v(\theta_n) \geq y \text{ then } \varphi_n(\theta_1, ..., \theta_n) = 1, \text{ and otherwise } \varphi_0(\theta_1, ..., \theta_n) = 1.$$

Let $\bar{s}(y)$ denote the strict threshold strategy with threshold y.

Let $\beta \in \mathcal{B}_{fosd}$ and $(c, \delta) \in C_{u_0}$. Let $U^P(\beta, c, \delta)$, $U^T(\beta, c, \delta)$, and $U^{ST}(\beta, c, \delta)$ be the maximal expected payoffs when the searcher is restricted to partition strategies, threshold strategies, and strict threshold strategies, respectively. Strict threshold strategies are more restrictive than threshold strategies, which in turn are more restrictive than partition strategies, so

$$U^{ST}(\beta, c, \delta) \leq U^T(\beta, c, \delta) \leq U^P(\beta, c, \delta). \tag{12}$$

Next, consider a hypothetical situation where the searcher learns the state ω after the very first interview. We will refer to this searcher as the *oracle*. Let $\hat{U}^*(\beta, c, \delta)$ be the maximal expected payoff of the oracle under prior β from the perspective of round 0.

Remark 2. *Because the oracle knows the state ω after the first draw, the optimal strategy for the oracle in state ω is Weitzman's (1979) solution conditional on the distribution of values F_ω. This solution is described by a weakly decreasing sequence of thresholds $(\tau^*_{\omega,t})_{t\in\mathbb{N}}$, where for each $t \in \mathbb{N}$, threshold $\tau^*_{\omega,t}$ is the unique solution of the equation*

$$\tau^*_{\omega,t} = \delta \left(-c + \int_0^\infty \max\{\tau^*_{\omega,t}, z\} F_\omega(\mathrm{d}z|x_{t+1}) \right). \tag{13}$$

Weitzman's (1979) condition that the sequence of thresholds must be weakly decreasing in t is satisfied under our assumption that $F_\omega(\cdot|x_t)$ is decreasing in t according to first order stochastic dominance.

Recall that $U^*(\beta, c, \delta)$ denotes the searcher's maximal expected payoff when choosing among all strategies. Observe that the searcher can only be better off if she learns the distribution she faces after the very first interview, so

$$U^*(\beta, c, \delta) \leq \hat{U}^*(\beta, c, \delta). \tag{14}$$

Note that inequality (14) is tight, in the sense that for each prior β there exists a sequence of priors $(\beta_k)_{k\in\mathbb{N}}$ such that

$$\lim_{k\to\infty} \beta_k = \beta \quad \text{and} \quad \lim_{k\to\infty} U^*(\beta_k) = \hat{U}^*(\beta, c, \delta).$$

This is because for each $\omega \in \Omega$, distribution λ_ω can be arbitrarily closely approximated by a sequence of priors $(\beta_k)_{k\in\mathbb{N}}$ that converges to β and has the following property. Choose each prior $\beta_k = (\lambda_{\omega,k}, p_\omega)_{\omega\in\Omega}$ such that for each $\omega', \omega'' \in \Omega$ and each $k \in \mathbb{N}$, distributions $\lambda_{\omega',k}$ and $\lambda_{\omega'',k}$ have disjoint support. In this case, the first interview reveals ω, so $U^*(\beta_k) = \hat{U}(\beta_k)$ for all $k \in \mathbb{N}$.

Summarizing the above, we obtain

$$u_0 \leq U^{ST}(\beta, c, \delta) \leq U^T(\beta, c, \delta) \leq U^P(\beta, c, \delta) \leq U^*(\beta, c, \delta) \leq \hat{U}^*(\beta, c, \delta). \tag{15}$$

The first inequality is because the searcher can guarantee u_0 by stopping in round 0. The second and third inequalities are by (12). The fourth inequality is because the highest attainable payoff when restricting to partition strategies is weakly smaller than that when using any strategies. The fifth inequality is by (14).

By (15), and because inequality (14) is tight, to prove the theorem it suffices to show that

$$\frac{\hat{U}^*(\beta, c, \delta)}{U^{ST}(\beta, c, \delta)} \leq 2 \quad \text{for all } \beta \in \mathcal{B}_{fosd} \text{ and all } (c, \delta) \in C_{u_0}, \tag{16}$$

and there exist $(c, \delta) \in C_{u_0}$ and a sequence of priors $(\beta_k)_{k\in\mathbb{N}}$ in $\mathcal{B}_{fosd}$ such that

$$\lim_{k\to\infty} \frac{\hat{U}^*(\beta_k, c, \delta)}{U^{ST}(\beta_k, c, \delta)} = \lim_{k\to\infty} \frac{\hat{U}^*(\beta_k, c, \delta)}{U^T(\beta_k, c, \delta)} = \lim_{k\to\infty} \frac{\hat{U}^*(\beta_k, c, \delta)}{U^P(\beta_k, c, \delta)} = 2. \tag{17}$$

In what follows, w.l.o.g. we assume that the oracle's optimal strategy prescribes to search at least one round, so

$$\hat{U}^*(\beta, c, \delta) > u_0,$$

This is because otherwise, if $\hat{U}^*(\beta, c, \delta) \leq u_0$, then the oracle prefers to stop immediately. But then so does the searcher who is constrained to strict threshold strategies, by using the threshold $y = 0$. In this case, inequality (15) holds as equality, in particular, we obtain $U^*(\beta, c, \delta)/U^{ST}(\beta, c, \delta) = 1 < 2$.

The proof of (16) and (17) is divided into four steps. Steps 1 and 2 simplify the problem by showing that it is sufficient to consider sets of priors that become simpler with each step. Inequality (16) is proved in Step 3. Equality (17) is proved in Step 4.

Step 1 shows that w.l.o.g. we can restrict attention to the priors such that the searcher's optimal strict threshold strategy has the threshold equal to the lower bound on the support.

Step 1. For each $\beta \in \mathcal{B}_{fosd}$ there exists $\hat{\beta} \in \mathcal{B}_{fosd}$ such that

$$\hat{U}^*(\beta, c, \delta) \leq \hat{U}^*(\hat{\beta}, c, \delta) \ \text{ and } \ U^{ST}(\beta, c, \delta) = U^{ST}(\hat{\beta}) = U_{\bar{s}(y)}(\hat{\beta}, c, \delta), \tag{18}$$

where $y = \inf\{v(\theta) : \theta \in \Theta_{\hat{\beta}}\}$.

Proof of Step 1. Let $\beta = (\lambda_\omega, p_\omega)_{\omega \in \Omega} \in \mathcal{B}_{fosd}$. Given a number $y > 0$, let $\bar{F}_{\omega,y}$ be the distribution with the floor y obtained from F_ω by pooling at y all the values below y. Formally, for each $\omega \in \Omega$ and each $i = 1, ..., n$ let

$$\bar{F}_{\omega,y}(\hat{v}|x_i) = F_\omega(\max\{\hat{v}, y\}|x_i), \quad \hat{v} \geq 0.$$

Let $\bar{\beta}_y = (\bar{F}_{\omega,y}, p_\omega)_{\omega \in \Omega}$. Observe that

$$\hat{U}^*(\beta, c, \delta) = \hat{U}^*(\bar{\beta}_0, c, \delta) \leq \hat{U}^*(\bar{\beta}_y, c, \delta), \tag{19}$$

where the equality is because $\bar{\beta}_0 = \beta$ and the inequality is because the oracle can only be better off as y increases.

Next, recall that a strict threshold strategy never chooses a candidate below its threshold, and stops immediately when the candidate is weakly above the threshold. So, as the lower bound on the support changes, it has no consequence on the payoff when this bound is below the threshold, and it improves the payoff when this bound is above the threshold. Thus we obtain that when the searcher uses a strict threshold strategy with a threshold $\tau > 0$, her expected payoff $U_{\bar{s}(\tau)}(\bar{\beta}_y, c, \delta)$ satisfies

$$U_{\bar{s}(\tau)}(\bar{\beta}_y, c, \delta) \text{ is } \begin{cases} \text{constant in } y \text{ for } y < \tau, \\ \text{continuously increasing in } y \text{ for } y \geq \tau. \end{cases} \tag{20}$$

Next, let z be the solution of the equation

$$\delta(-c+z) = U^{ST}(\beta, c, \delta). \tag{21}$$

Let $V(y)$ be the expected payoff of the searcher who uses the strict threshold strategy whose threshold is equal to the lower bound on the support, so

$$V(y) = U_{\bar{s}(y)}(\bar{\beta}_y, c, \delta).$$

We have

$$V(0) = U_{\bar{s}(0)}(\beta, c, \delta) \le \sup_{\tau \ge 0} U_{\bar{s}(\tau)}(\beta, c, \delta) = U^{ST}(\beta, c, \delta) \le V(z), \tag{22}$$

where the first equality is because $\bar{\beta}_0 = \beta$, the second equality is by the definition of $U^{ST}(\beta, c, \delta)$, and the second inequality is because when the lower bound on the support is z, the expected payoff cannot be smaller than $\delta(-c+z)$, which is equal to $U^{ST}(\beta, c, \delta)$ by (21).

Next, by (20) and (22), there exists $y \in [0, z]$ such that

$$V(y) = U^{ST}(\beta, c, \delta) = \sup_{\tau \ge 0} U_{\bar{s}(\tau)}(\bar{\beta}_y, c, \delta) = U^{ST}(\bar{\beta}_y, c, \delta). \tag{23}$$

In particular, y is the optimal threshold under $\bar{\beta}_y$.

In summary, by (19) and (23), we obtain (18) with $\hat{\beta} = \bar{\beta}_y$. This completes the proof of Step 1. □

Step 2 shows that w.l.o.g. we can restrict attention to the priors that have support on at most two values.

Step 2. For each $\beta \in \mathcal{B}_{fosd}$ there exist $\hat{\beta} \in \mathcal{B}_{fosd}$ and $y, z \in \mathbb{R}_+$ such that $y \le z$ and

$$\begin{gathered} \Theta_{\hat{\beta}} = \{\theta', \theta''\} \text{ with } v(\theta') = y \text{ and } v(\theta'') = z, \\ \hat{U}^*(\beta, c, \delta) \le \hat{U}^*(\hat{\beta}, c, \delta), \quad \text{and} \quad U^{ST}(\beta, c, \delta) = U^{ST}(\hat{\beta}, c, \delta). \end{gathered} \tag{24}$$

Proof of Step 2. Let $\beta = (F_\omega, p_\omega)_{\omega \in \Omega} \in \mathcal{B}_{fosd}$. Let y be the lowest nonnegative value and let z be the highest value of the types in Θ_β, so

$$y = \max\{0, \inf\{v(\theta) : \theta \in \Theta_\beta\}\} \quad \text{and} \quad z = \sup\{v(\theta) : \theta \in \Theta_\beta\} = v(\bar{\theta}).$$

In what follows we consider $y < z$, as otherwise (24) trivially holds with $\hat{\beta} = \beta$.

For each $\omega \in \Omega$ and each $t = 1, ..., n$ let $q_{\omega,t}$ be the probability such that

$$(1 - q_{\omega,t})y + q_{\omega,t}z = \int_0^\infty \hat{v} \mathrm{d}F_\omega(\hat{v}|x_t).$$

Given $x \in [y, z]$, let $\xi_{\hat{v}}$ be the probability of z under the mean-preserving lottery over $\{y, z\}$ with mean $\hat{v}$, so

$$(1 - \xi_{\hat{v}})y + \xi_{\hat{v}}z = \hat{v}. \tag{25}$$

Define $\hat{\beta} = (\hat{F}_\omega, p_\omega)_{\omega\in\Omega}$ as follows. For each $\omega \in \Omega$ and each $t = 1, ..., n$ let $\hat{F}_\omega(\cdot|x_t)$ be obtained from $F_\omega(\cdot|x_t)$ by replacing the value $v(\theta)$ of each draw of $\theta \in \Theta_\beta$ with the mean-preserving lottery over y and z with probabilities $1 - \xi_\theta$ and ξ_θ, respectively. After integrating over all possible draws of θ from $F_\omega(\cdot|x_t)$, the resulting distribution over y and z assigns probability $1 - q_{\omega,t}$ to y and $q_{\omega,t}$ to z. In particular, $\Theta_{\hat{\beta}} = \{\theta', \theta''\}$ with $v(\theta') = y$ and $v(\theta'') = z$.

First let us show that $\hat{U}^*(\hat{\beta}, c, \delta) \geq \hat{U}^*(\beta, c, \delta)$. Let $u^*_{\omega,t}(\hat{v})$ be the maximal expected payoff of the oracle (excluding u_0) from perspective of round t (after having interviewed candidate t) when the best interviewed candidate up to round t has value $\hat{v}$. It suffices to show that

$$(1 - \xi_{v(\theta)})u^*_{\omega,t}(y) + \xi_{v(\theta)}u^*_{\omega,t}(z) \geq u^*_{\omega,t}(v(\theta)) \tag{26}$$

for each $\theta \in \Theta_\beta$, each $t = 1, ..., n$, and each $\omega \in \Omega$.

Let $\tau^*_{\omega,t}$ be the oracle's optimal threshold under F_ω in round t as given by (13). Because the oracle stops and obtains $\hat{v}$ whenever $\hat{v} \geq \tau^*_{\omega,t}$, and prefers to continue rather than stopping and having $\hat{v}$ whenever $\hat{v} < \tau^*_{\omega,t}$, we have

$$u^*_{\omega,t}(\hat{v}) = \hat{v} \text{ if } \hat{v} \geq \tau^*_{\omega,t}, \quad \text{and} \quad u^*_{\omega,t}(\hat{v}) \geq \hat{v} \text{ if } \hat{v} < \tau^*_{\omega,t}.$$

Moreover, as follows from Weitzman (1979) and the Envelope Theorem, the marginal increment of $u^*_{\omega,t}(\hat{v})$ is increasing in $\hat{v}$, so $u^*_{\omega,t}(\hat{v})$ is weakly convex in $\hat{v}$. Consequently, we obtain (26).

Second, we show that

$$U^{ST}(\beta, c, \delta) = U_{\bar{s}(y)}(\beta, c, \delta) = U_{\bar{s}(y)}(\hat{\beta}, c, \delta) = U^{ST}(\hat{\beta}, c, \delta).$$

Let μ_t be the expected value of candidate t under β. Note that this value is the same under $\hat{\beta}$ by the construction of $\hat{\beta}$, and it is given by

$$\mu_t = \sum_{\omega\in\Omega}\left(\int_0^\infty \hat{v}\mathrm{d}F_\omega(\hat{v}|x_t)\right)p_\omega = \sum_{\omega\in\Omega}\left((1-q_{\omega,t})y + q_{\omega,t}z\right)p_\omega. \tag{27}$$

By Step 1, w.l.o.g. we assume that under β the searcher's optimal threshold is the lower-bound value y, so

$$U^{ST}(\beta,c,\delta) = U_{\bar{s}(y)}(\beta,c,\delta).$$

Clearly, when the searcher uses the strategy with threshold y, she stops in round 0 or 1, so her expected payoff is the same under β and $\hat{\beta}$, and it is given by

$$U_{\bar{s}(y)}(\beta,c,\delta) = U_{\bar{s}(y)}(\hat{\beta},c,\delta) = u_0 + \max\{0,\delta(-c+\mu_1)\}.$$

It remains to show that threshold y is optimal under $\hat{\beta}$, and therefore $U_{\bar{s}(y)}(\hat{\beta},c,\delta) = U^{ST}(\hat{\beta},c,\delta)$.

By assumption, $F_\omega(\cdot|x_t)$ is decreasing in t according to f.o.s.d. That is, conditional on any given threshold τ, the searcher's expected payoff can only get worse with time. Consequently, as y is the optimal threshold from perspective of round 0, it follows that y must remain the optimal threshold from perspective of every round $t=1,...,n-1$. This means that when the value of the best interviewed candidate up to round t is y, then the searcher prefers to stop rather than to search one more round and then stop,

$$y \geq \delta(-c+\mu_t) \ \text{ for each } t=2,...,n, \tag{28}$$

where μ_t is the expected value in round t given by (27).

Let $\underline{u}_t^{ST}(\tau,\beta)$ be the expected payoff (excluding u_0) from the perspective of round t when the searcher uses the strict threshold strategy with threshold τ, and the value of the best interviewed candidate up to round t is equal to y. We show by induction that for every threshold $\tau > y$ we have

$$\underline{u}_1^{ST}(\tau,\hat{\beta}) \leq y. \tag{29}$$

This inequality explains the searcher's optimal choice in round 1 when she drew the worst possible value, y. The right-hand side of (29) is the payoff of stopping in round 1 and getting that value. The left-hand side of (29) is the expected payoff when continuing the search with threshold τ. Inequality (29) means stopping is preferred. This proves that threshold τ is weakly inferior to threshold y.

We now provide the induction argument to prove (29). In the very last round n, if the value of the best candidate interviewed so far is y, the expected payoff in this round cannot exceed y, so

$$\underline{u}_n^{ST}(\tau, \hat{\beta}) \leq y.$$

We proceed by induction in $t = n-1, n-2, ..., 2$. Suppose that

$$\underline{u}_t^{ST}(\tau, \hat{\beta}) \leq y. \tag{30}$$

Then

$$\begin{aligned}
\underline{u}_{t-1}^{ST}(\tau, \hat{\beta}) &= \delta \left(-c + \left(\sum_{\omega \in \Omega} q_{\omega,t} p_\omega \right) z + \left(1 - \sum_{\omega \in \Omega} q_{\omega,t} p_\omega \right) \underline{u}_t^{ST}(\tau, \hat{\beta}) \right) \\
&\leq \delta \left(-c + \left(\sum_{\omega \in \Omega} q_{\omega,t} p_\omega \right) z + \left(1 - \sum_{\omega \in \Omega} q_{\omega,t} p_\omega \right) y \right) \\
&= \delta \left(-c + \mu_t \right) \leq y.
\end{aligned}$$

The first line is because when the best value up to $t-1$ is y, which is smaller than the threshold τ, the search must proceed to round t. In round t, the searcher gets z with probability $q_{\omega,t}$ in each state ω. With the complementary probability she gets y, in which case her continuation payoff is $\underline{u}_t^{ST}(\tau, \hat{\beta})$. The second line is by the induction assumption (30). The third line is by (27) and (28). This completes the induction argument and the proof of Step 2. □

Step 3 shows inequality (16).

Step 3. For each $\beta \in \mathcal{B}_{fosd}$, $\hat{U}^*(\beta, c, \delta) \leq 2U^{ST}(\beta, c, \delta)$.

Proof of Step 3. Let $\beta \in \mathcal{B}_{fosd} = (F_\omega, p_\omega)_{\omega \in \Omega}$. By Step 2, w.l.o.g., suppose that $\Theta_\beta = \{\theta', \theta''\}$ with $v(\theta') = y$, $v(\theta'') = z$, and $y < z$. We can exclude the case of $y = z$, because in that case β is degenerate, and $\hat{U}^*(\beta, c, \delta) = U^{ST}(\beta, c, \delta)$.

We represent each distribution $F_\omega(\cdot|x_t)$ as the lottery over y and z with probabilities $1 - q_{\omega,t}$ and $q_{\omega,t}$, respectively. Thus, each state ω induces a sequence of probabilities $(q_{\omega,1}, ..., q_{\omega,n})$. Because we have assumed that $F_\omega(\cdot|x_t)$ is ordered in t according to f.o.s.d., this sequence of probabilities is decreasing, so

$$q_{\omega,1} \geq q_{\omega,2} \geq ... \geq q_{\omega,n} \text{ for each } \omega \in \Omega.$$

Let

$$Q_{t,\Omega} = (1 - q_{1,\omega})(1 - q_{2,\omega})...(1 - q_{t,\omega}).$$

Consider the oracle's optimal payoff. Because we have only two values in the support, y and z, it follows from Weitzman (1979) that the oracle's optimal strategy is described by a cutoff round $T_\omega \in \{1, ..., n\}$ such that the oracle stops the search when z is drawn or when round T_ω is reached. The oracle's optimal payoff is given by

$$\hat{U}^*(\beta, c, \delta) = \sum_{\omega \in \Omega} v^*_\omega p_\omega, \tag{31}$$

where v^*_ω is the oracle's expected payoff in state ω,

$$\begin{aligned} v^*_\omega = u_0 + \delta(-c + q_{1,\omega} z) + \delta^2 Q_{1,\omega}(-c + q_{2,\omega} z) + ... + \delta^{T_\omega} Q_{T_\omega - 1,\omega}(-c + q_{T_\omega,\omega} z) \\ + \delta^{T_\omega} Q_{T_\omega,\omega} y. \end{aligned} \tag{32}$$

W.l.o.g. suppose that prior β satisfies the following property:

$$q_{t,\omega} = 0 \ \text{ for each } t \in \{T_\omega + 1, ..., n\} \text{ and each } \omega \in \Omega. \tag{33}$$

In words, it is impossible to draw z after the round when the oracle stops the search. Changing a positive $q_{t,\omega}$ to zero in these rounds has no effect on the oracle's payoff, but it weakly reduces the payoff of the searcher who uses a strict threshold strategy. So, overall, the ratio of the oracle's payoff to the optimal threshold payoff increases.

Next, consider the payoffs under strict threshold strategies. Because we have only two values in the support, y and z, we only need to consider two strict thresholds strategies: one with threshold y that stops in round 1, and the other with threshold z that continues searching until z is drawn or the entire pool of candidates is interviewed. Let v^y_ω and v^z_ω be the searcher's payoffs in state ω when using thresholds y and z, respectively. We have

$$v^y_\omega = u_0 + \delta(-c + q_{1,\omega} z + (1 - q_{1,\omega}) y) = \delta(-c + q_{1,\omega} z) + \delta Q_{1,\omega} y. \tag{34}$$

We also have

$$\begin{aligned} v^z_\omega &= u_0 + \delta(-c + q_{1,\omega} z) + \delta^2 Q_{1,\omega}(-c + q_{2,\omega} z) + ... + \delta^{T_\omega} Q_{T_\omega - 1,\omega}(-c + q_{T_\omega,\omega} z) \\ &\quad - \delta^{T_\omega} Q_{T_\omega,\omega}(\delta + \delta^2 + ... + \delta^{n - T_\omega}) c \\ &= v^*_\omega - \delta^{T_\omega} Q_{T_\omega,\omega} \left(y + \frac{\delta(1 - \delta^{n - T_\omega})}{1 - \delta} c \right), \end{aligned} \tag{35}$$

where we used (33) to establish the first equality, and the second equality is by the definition of v^*_ω in (32). The expected payoffs from the two threshold strategies are given by

$$U_{\bar{s}(y)}(\beta, c, \delta) = \sum_{\omega \in \Omega} v^y_\omega p_\omega \quad \text{and} \quad U_{\bar{s}(z)}(\beta, c, \delta) = \sum_{\omega \in \Omega} v^z_\omega p_\omega. \tag{36}$$

Next, we have

$$\begin{aligned} v^*_\omega - v^z_\omega - v^y_\omega &= \delta^{T_\omega} Q_{T_\omega,\omega} \left(y + \frac{\delta(1-\delta^{n-T_\omega})}{1-\delta} c \right) - \delta(-c + q_{1,\omega} z) - \delta Q_{1,\omega} y - u_0 \\ &= \left(\delta^{T_\omega - 1} Q_{T_\omega,\omega} - Q_{1,\omega} \right) \delta y + \left(\delta^{T_\omega} Q_{T_\omega,\omega} \frac{(1-\delta^{n-T_\omega})}{1-\delta} + 1 \right) \delta c - \delta q_{1,\omega} z - u_0 \\ &\le \frac{1-\delta^n}{1-\delta} \delta c - u_0. \end{aligned} \tag{37}$$

The first equality is by (32), (34), and (35). The second equality is by rearranging of terms. The first inequality is because $Q_{T_\omega,\omega} \le Q_{1,\omega} \le 1$, so $\delta^{T_\omega - 1} Q_{T_\omega,\omega} - Q_{1,\omega} \le 0$ and

$$\begin{aligned} \delta^{T_\omega} Q_{T_\omega,\omega} \frac{(1-\delta^{n-T_\omega})}{1-\delta} + 1 \le \delta^{T_\omega} \frac{(1-\delta^{n-T_\omega})}{1-\delta} + 1 &= \left(\delta^{T_\omega} + \delta^{T_\omega+1} + \ldots + \delta^{n-1} \right) + 1 \\ &\le 1 + \delta + \ldots + \delta^{n-1} = \frac{1-\delta^n}{1-\delta}. \end{aligned}$$

By (31), (36), and (37) we obtain

$$\begin{aligned} \hat{U}^*(\beta, c, \delta) - U_{\bar{s}(z)}(\beta, c, \delta) - U_{\bar{s}(y)}(\beta, c, \delta) &= \sum_{\omega \in \Omega} (v^*_\omega - v^z_\omega - v^y_\omega) p_\omega \\ &\le \frac{1-\delta^n}{1-\delta} \delta c - u_0 \le 0, \end{aligned} \tag{38}$$

where the last inequality is by assumption (1).

Finally, the searcher's payoff from the optimal threshold strategy is given by

$$U^{ST}(\beta, c, \delta) = \max\{U_{\bar{s}(y)}(\beta, c, \delta), U_{\bar{s}(z)}(\beta, c, \delta)\}. \tag{39}$$

Thus, we obtain

$$\begin{aligned} \frac{\hat{U}^*(\beta, c, \delta)}{U^{ST}(\beta, c, \delta)} &= \frac{U_{\bar{s}(z)}(\beta, c, \delta)}{U^{ST}(\beta, c, \delta)} + \frac{\hat{U}^*(\beta, c, \delta) - U_{\bar{s}(z)}(\beta, c, \delta)}{U^{ST}(\beta, c, \delta)} \\ &\le \frac{U_{\bar{s}(z)}(\beta, c, \delta)}{U_{\bar{s}(z)}(\beta, c, \delta)} + \frac{\hat{U}^*(\beta, c, \delta) - U_{\bar{s}(z)}(\beta, c, \delta)}{U_{\bar{s}(y)}(\beta, c, \delta)} \le 2, \end{aligned}$$

where the first inequality is by (39), and the second inequality is by (38). This completes the proof of Step 3. □

Step 4 proves the tightness of the bound.

Step 4. There exist $(c, \delta) \in C_{u_0}$ and a sequence of priors $(\beta_k)_{k\in\mathbb{N}}$ in $\mathcal{B}_{fosd}$ such that

$$\lim_{k\to\infty} \frac{\hat{U}^*(\beta_k, c, \delta)}{U^{ST}(\beta_k, c, \delta)} = \lim_{k\to\infty} \frac{\hat{U}^*(\beta_k, c, \delta)}{U^{T}(\beta_k, c, \delta)} = \lim_{k\to\infty} \frac{\hat{U}^*(\beta_k, c, \delta)}{U^{P}(\beta_k, c, \delta)} = 2. \tag{40}$$

Proof of Step 4. Given $u_0 > 0$, consider $(c, \delta) \in C_{u_0}$ that satisfy

$$(\delta^2 + \delta^3 + ... + \delta^n)c = u_0. \tag{41}$$

Let $(p_k, q_k, z_k)_{k\in\mathbb{N}}$ be a sequence such that

$$p_k, q_k \in (0, 1] \text{ for each } k \in \mathbb{N},\ \lim_{k\to\infty} p_k = 0, \text{ and } \lim_{k\to\infty} q_k = 0,$$

and

$$z_k = \frac{c}{p_k q_k}.$$

Using the above definition of z_k, we obtain the next expression that will be used later,

$$p_k\delta(-c + q_k z_k) = (1 - p_k)\delta c. \tag{42}$$

Consider a sequence of priors $(\beta_k)_{k\in\mathbb{N}}$, where each β_k is described as follows. There are two states, 0 and 1, that are realized with probabilities $1-p_k$ and p_k, respectively. When state 0 is realized, then all candidates have value 0 with certainty. When state 1 is realized, then the value of each candidate is 0 with probability $1 - q_k$ and z_k with probability q_k, independently of other candidates. Note that β_k does not satisfy our assumption of the common support across the states. This simplifies the exposition. The proof can be easily adjusted to accommodate a probability of value z_k in state 0 that is positive but fast vanishing as $k \to \infty$.

Observe that under each β_k, as there are only two types in Θ_β, threshold and partition strategies coincide. Moreover, because the lower type has value 0, so it is as good as the outside option, the payoffs of threshold and strict threshold strategies are equal. We thus obtain

$$U^{ST}(\beta_k, c, \delta) = U^{T}(\beta_k, c, \delta) = U^{P}(\beta_k, c, \delta). \tag{43}$$

Next, we compute the payoff from the optimal strict threshold strategy, $U^{ST}(\beta_k, c, \delta)$. We only need to compare the outside option whose payoff is

$$U_{\bar{s}(0)}(\beta_k, c, \delta) = u_0. \tag{44}$$

and the strategies with two thresholds, y and z_k, for an arbitrary $y \in (0, z_k)$. Threshold y means stopping after the first interview, so

$$U_{\bar{s}(y)}(\beta_k, c, \delta) = u_0 + (1 - p_k)\delta(-c) + p_k\delta(-c + q_k z_k) = u_0, \tag{45}$$

where the first equality is by the definition of β_k and the second equality is by (42). Threshold z_k means searching until z_k is drawn, or until all the candidates are interviewed, so

$$\begin{aligned} U_{\bar{s}(z)}(\beta_k, c, \delta) &= u_0 + (1 - p_k)\delta(-c)(1 + \delta + ... + \delta^{n-1}) \\ &\quad + p_k\delta(-c + q_k z_k)(1 + (1 - q_k)\delta + ... + (1 - q_k)^{n-1}\delta^{n-1}) \\ &= u_0 + (1 - p_k)\delta c\big(\delta((1 - q_k) - 1) + ... + \delta^{n-1}((1 - q_k)^{n-1} - 1)\big) \\ &\leq u_0, \end{aligned} \tag{46}$$

where the first equality is by the definition of β_k, the second equality is by (42), and the inequality is because $q_k \geq 0$. Consequently, by (44), (45), and (46), we obtain

$$U^{ST}(\beta_k, c, \delta) = \max\left\{U_{\bar{s}(0)}(\beta_k, c, \delta), U_{\bar{s}(y)}(\beta_k, c, \delta), U_{\bar{s}(z)}(\beta_k, c, \delta)\right\} = u_0. \tag{47}$$

Next, we find the payoff from the oracle's optimal strategy, $\hat{U}^*(\beta_k, c, \delta)$. The oracle stops in round 1 when the state is 0, and searches until z_k is drawn or until all the candidates are interviewed when the state is 1. Thus,

$$\begin{aligned} \hat{U}^*(\beta_k, c, \delta) &= u_0 + (1 - p_k)\delta(-c) \\ &\quad + p_k\delta(-c + q_k z_k)(1 + (1 - q_k)\delta + ... + (1 - q_k)^{n-1}\delta^{n-1}) \\ &= u_0 + (1 - p_k)\delta c\big((1 - q_k)\delta + ... + (1 - q_k)^{n-1}\delta^{n-1}\big), \end{aligned}$$

where the first equality is by the definition of β_k and the second equality is by (42). Taking the limit as $k \to \infty$, so $p_k \to 0$ and $q_k \to 0$, we obtain

$$\lim_{k\to\infty} \hat{U}^*(\beta_k, c, \delta) = u_0 + \delta c(\delta + ... + \delta^{n-1}) = 2u_0, \tag{48}$$

where the second equality is by (41). Consequently, by (43), (47) and (48) we obtain (40). This completes the proof of Step 4. □

References

Arrow, K. J., A. Sen, and K. Suzumura (2010): *Handbook of social choice and welfare*, vol. 2, Elsevier.

Babaioff, M., N. Immorlica, D. Kempe, and R. Kleinberg (2008): "Online auctions and generalized secretary problems," *ACM SIGecom Exchanges*, 7, 1–11.

Barocas, S. and A. D. Selbst (2016): "Big Data's disparate impact," *California Law Review*, 104, 671–732.

Becker, G. (1957): *The Economics of Discrimination*, University of Chicago Press.

Berk, R. (2012): *Criminal justice forecasts of risk: A machine learning approach*, Springer Science & Business Media.

Berk, R. A. and J. Bleich (2013): "Statistical procedures for forecasting criminal behavior: A comparative assessment," *Criminology & Public Policy*, 12, 513.

Bertrand, M. and E. Duflo (2017): "Field Experiments on Discrimination," in *Handbook of Economic Field Experiments*, ed. by A. Banerjee and E. Duflo, North-Holland, vol. 1, 309–393.

Brennan, T. and W. L. Oliver (2013): "Emergence of machine learning techniques in criminology: Implications of complexity in our data and in research questions," *Criminology & Public Policy*, 12, 551.

Burdett, K. and K. L. Judd (1983): "Equilibrium price dispersion," *Econometrica*, 51, 955–969.

Chouldechova, A. (2017): "Fair prediction with disparate impact: A study of bias in recidivism prediction instruments," *Big Data*, 5, 153–163.

Colquitt, J. A. and J. B. Rodell (2015): "Measuring Justice and Fairness," in *The Oxford Handbook of Justice in the Workplace*, ed. by R. Cropanzano and M. L. Ambrose, Oxford University Press.

Corbett-Davies, S. and S. Goel (2018): "The measure and mismeasure of fairness: A critical review of fair machine learning," ArXiv:1808.00023.

DeGroot, M. H. (1968): "Some problems of optimal stopping," *Journal of the Royal Statistical Society: Series B (Methodological)*, 30, 108–122.

Dwork, C. and C. Ilvento (2018): "Individual fairness under composition," *Proceedings of Fairness, Accountability, Transparency in Machine Learning*.

Feldman, M., S. A. Friedler, J. Moeller, C. Scheidegger, and S. Venkatasubramanian (2015): "Certifying and removing disparate impact," in *proceedings of the 21th ACM SIGKDD international conference on knowledge discovery and data mining*, 259–268.

Ferguson, T. S. (1989): "Who Solved the Secretary Problem?" *Statistical Science*, 4, 282–289.

Fox, J. H. and L. G. Marnie (1960): "In Martin Gardner's column: Mathematical games," *Scientific American*, 202, 150–153.

Freeman, P. R. (1983): "The Secretary Problem and Its Extensions: A Review," *International Statistical Review*, 51, 189–206.

Gastwirth, J. L. (1976): "On probabilistic models of consumer search for information," *Quarterly Journal of Economics*, 90, 38–50.

Kleinberg, J., S. Mullainathan, and M. Raghavan (2016): "Inherent trade-offs in the fair determination of risk scores," ArXiv:1609.05807.

Krengel, U. and L. Sucheston (1978): "On semiamarts, amarts, and processes with finite value," in *Probability on Banach spaces*, Dekker, New York, vol. 4 of *Advances in Probability and Related Topics*, 197–266.

McCall, J. J. (1970): "Economics of information and job search," *Quarterly Journal of Economics*, 84, 113–126.

Neumark, D. (2018): "Experimental research on labor market discrimination," *Journal of Economic Literature*, 56, 799–866.

Parakhonyak, A. and A. Sobolev (2015): "Non-Reservation Price Equilibrium and Search without Priors," *Economic Journal*, 125, 887–909.

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

Rothschild, M. (1974): "Searching for the Lowest Price When the Distribution of Prices is Unknown," *Journal of Political Economy*, 82, 689–711.

Schlag, K. and A. Zapechelnyuk (2021): "Robust sequential search," *Theoretical Economics*, 16, 1431–1470.

Sleator, D. and R. Tarjan (1985): "Amortized efficiency of list update and paging rules," *Communications of the ACM*, 28, 202–208.

Stigler, G. J. (1962): "Information in the labor market," *Journal of Political Economy*, 70, 94–105.

Vanderschraaf, P. (2023): *Bargaining Theory*, Cambridge University Press.

Weitzman, M. L. (1979): "Optimal Search for the Best Alternative," *Econometrica*, 47, 641–654.