FAIR HIRING PROCEDURES

Karl H. Schlag and Andriy Zapechelnyuk

Abstract. While much is known about statistical and behavioral discrimination in hiring decisions, little attention has been given to how the organization of the interviewing process affects fairness. Using a sequential search model, we define fairness through two principles—equal treatment of equals and invariance under reordering—and fully characterize the hiring procedures that satisfy them. We show that such procedures have a simple structure and can be preregistered with a regulator to ensure fair treatment of applicants. Our analysis reveals that several common hiring practices are unfair according to these principles of fairness.

JEL Classification: D81, D83, J71

Keywords: Sequential search, hiring procedure, fairness, discrimination, equal opportunity, shortlist, categorization

September 25, 2025

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, Leeat Yariv, and participants of various seminars where this paper has been presented.

1. Introduction

Fair employment is defined as "Employment of workers on a basis of equality without discrimination or segregation especially because of race, color, or creed." (Merriam-Webster Dictionary). Numerous studies investigate discrimination and statistical biases in hiring that may arise when assessing applicants, as surveyed in Bertrand and Duflo (2017) and Neumark (2018). However, there is little research on how the process of interviewing and selecting applicants influences fairness. The aim of this paper is to fill this gap. We formalize what it means for a hiring process to be fair, and characterize hiring procedures that are fair. Our findings reveal that unfair treatment is intrinsic in theoretically optimal procedures, as well as in many of those observed in practice. We provide novel guidelines (or best practices) for how to organize the hiring process to make it fair.

We investigate fairness of a hiring process within the following model of sequential search. There is a pool of applicants for a job vacancy, which comprises of all who applied for a job and fit the job description. Each applicant has an identity that incorporates all the publicly observable information about the applicant, such as their curriculum vitae or resume. Each applicant also has a job fit that includes the information relevant for the employer's hiring decision, such as professionalism, reliability, ability to learn and adapt, communication skills, work ethic, etc. An applicant's job fit is only revealed when that applicant is interviewed. The employer observes the pool of applicants and decides in what order to interview them, when to stop interviewing, and whom to hire (if anyone). These choices of the employer are referred to as the hiring procedure.

The value of hiring an applicant is determined by their job fit. No discrimination means that only the job fit should matter for hiring. This leads us to two fairness principles. The first principle is called *equal treatment of equals*. When two applicants have the same job fit, regardless of whether either of them have been interviewed, they should not be discriminated. Hence, as they both cannot be hired for the same vacancy, we demand that they have the same chance to be hired. The second principle is called *invariance under reordering*. Applicants should not be discriminated by the order in which interviews have taken place. Hence, we demand that no applicant would be better off or worse off if some interviewed applicant was interviewed earlier.

We show that a procedure is fair if and only if it is a categorization procedure. According to such a procedure, first, the employer makes an irrevocable commitment to a categorization that divides the applicants' job fits into two categories, named *strong* and *weak*. Then, the hiring process starts. An interviewee is chosen at random, with equal probability. If the job fit of the interviewed applicant turns out to be strong, then this applicant is hired on the spot. If that job fit turns out to be weak, then the next interviewee is chosen at random, with equal probability, among the uninterviewed applicants, and the process continues as above. If all applicants have been interviewed but none of them is strong, then the employer can either choose one of the interviewed applicants or choose to hire nobody. Applicants who have the same job fit must be hired equally likely, and who is hired must not depend on the order in which the applicants were interviewed.

We obtain four implications of fair hiring. First of all, everyone must be given equal chance. Observable characteristics of applicants may not be used to differentiate them. Even if an observable characteristic contains some information about the job fit, it must be ignored and cannot influence hiring until the job fit is revealed during the interview. Second of all, the hiring criteria must be committed to prior to interviews. So, the employer is not allowed to change the categorization of the job fits once the interviews have started, regardless of what is discovered during the interviews. Third of all, we obtain a side effect that categorization procedures are easy to monitor. Transparency of the procedures simplifies investigation of complaints of unlawful discrimination and allows to challenge any wrongdoing of the employer. In fact, we advocate that hiring should start by preregistering categories with a regulator. Finally, fair procedures are simple to implement. This stands in contrast with optimal procedures without fairness constraints that are technically challenging and are characterized only in a few special cases (DeGroot, 1968; Rothschild, 1974; Weitzman, 1979).

The concern for fairness described in this paper has practical consequences. According to Lazear et al. (2018): "Being hired into a job depends not only on one's own skill but also on that of other applicants. When another able applicant applies, a well-suited worker may be forced into unemployment or into accepting an inferior job." Lazear et al. (2018) provide strong empirical evidence from the US labor market of this practice being commonplace. According to our findings, this hiring practice cannot be considered fair, as it prescribes the hiring criteria to be adapted to what an employer learns about applicants' job fits. In a similar vein, optimal solutions to some classic search models cannot be considered fair. Here we refer to interviewing applicants in the order according to how good their resumes are (as in the sequential search of Weitzman, 1979) and hiring only those that whose job fit is superior to an initial set of interviewed applicants (as in the secretary problem of Fox and Marnie, 1960). Similarly, with the help of our paper, we identify that some common hiring practices are unfair. This includes hiring in batches, interviewing in any fixed order, and terminating the search with no hire before all applicants in the pool have been interviewed.

As mentioned before, our result implies that the applicants' observable characteristics should play no role if fair hiring. However, in practice, observable characteristics may be used to preselect among applicants, especially when the employer operates under budget constraint and cannot afford interviewing all applicants. In an extension to our model (Section 4.2), we relax our first principle and only impose it on those who can be potentially interviewed. As a result, fair hiring permits shortlisting. Observable characteristics may be used to determine a shortlist—subset of applicants preselected to be interviewed—but may not play any further role after the shortlist is formed.

Related Literature. This paper is related to the 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 recommendations. This literature is surveyed in Bertrand and Duflo (2017) and Neumark (2018). Additionally, the phenomenon of statistical discrimination has recently received a surge of attention in the computer science literature, 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 to correct 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 et al., 2024). Unlike the literature mentioned above, our paper is not concerned with discrimination due to assessment biases. In fact, we assume that interviews reveal accurate information about the applicants. 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), the law and justice literature has no unified view on what fairness is. Moreover, this literature lacks formalism when dealing with the concept of fairness. In the economics literature, the concern for fairness formally appears in the bargaining and social choice context. Therein, 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, the corresponding concept is called distributive or allocative fairness (Kagel and Roth, 1995). This is an ex-post concept. In contrast, the computer science literature has been interested in ex-ante fairness. Randomization in procedures is used to achieve a statistical balance between different population groups (e.g., Kleinberg et al., 2017; Dwork and Ilvento, 2018; Corbett-Davies et al., 2024).[1] Neither of these two strands of literature is of any help to us. A hiring procedure cannot be equitable ex-post as everyone wants the job but only one applicant gets it. It also should not be made equitable ex-ante, as this would mean to hire a random applicant or to commit to interview all applicants in the pool before hiring. Instead, we capture the intermediate nature of procedural fairness, using the canonical model of sequential search due to Stigler (1962) and McCall (1970).

The paper is organized as follows. In Section 2 we introduce a model of hiring. In Section 3 we postulate two fairness principles, characterize fair procedures, discuss the implications, and provide the intuition and counterexamples. Two extensions of the model are presented in Section 4. Section 5 concludes. The proof of our result is in Appendix A.

[1] Mitchell et al. (2021) survey the literature on algorithmic fairness, and Bolton et al. (2005) provide an experimental evidence of the distinction between ex-ante and ex-post fairness.

2. Model

An employer wishes fill to a job vacancy. There is a pool of n applicants that consists of those who applied for the vacancy and fit the job description. Each applicant i has a profile θ_i of attributes, $i = 1, ..., n$. This profile is called the job fit. It includes all the information relevant for the employer's hiring decision, such as professionalism, reliability, ability to learn and adapt, communication skills, work ethic, etc. There is also an outside option, denoted by $i = 0$, with an associated profile θ_0. It captures the outcome if the employer does not hire anybody. To simplify the exposition, we will treat $i = 0$ as an additional applicant whose job fit θ_0 is known. For instance, this can be an existing employee whose contract came to an end, so the employer may extend their contract instead of hiring a new employee.

Let $N = \{0, 1, ..., n\}$. Let Θ be the set of possible job fits, so $\theta_i \in \Theta$ for all $i \in N$. Assume that Θ is finite. Job fits are partially ordered. The partial order $\succ$ will be used to evaluate fairness from the viewpoint of an outside observer. This allows both for a setting where there is a complete order, as well as a setting where it is hard to tradeoff different attributes. For example, a job fit could be a profile of L numeric attributes, so $\Theta \subset \mathbb{R}^L$. Then $\theta \succ \theta'$ if θ exceeds θ' in all attributes and strictly exceeds it in some attribute. We further assume that there exists a so-called *ideal* job fit in Θ, denoted by $\bar{\theta}$, that is better than all other job fits, so $\bar{\theta} \succ \theta$ for all $\theta \in \Theta \backslash \{\bar{\theta}\}$.

The employer is in charge of interviewing applicants and selecting which one to hire. Initially, the employer knows the job fit θ_0 of applicant 0, but not the job fits of any of the other applicants. To discover these job fits, the employer interviews applicants one by one.

The search for a suitable job applicant proceeds in rounds. In each round $t = 1, 2, ..., n$, the employer selects one of the uninterviewed applicants and conducts an interview to discover the job fit of that applicant. Upon discovering the job fit, the employer decides whether to stop the interviewing process and hire one of the interviewed applicants, or to proceed to the next round. In round n the process stops.

We refer to the way in which the employer navigates this interviewing process as a hiring procedure. Specifically, a hiring procedure prescribes who to interview next, when to stop interviewing, and whom to hire. It is given by a triple $\pi = (b, s, a)$

defined as follows. Let history h_t in round $t = 1, ..., n$ be the list of applicants interviewed up to round t together with their job fits (excluding applicant 0 whose job fit θ_0 is fixed), so $h_t = ((i_1, \theta_{i_1}), ..., (i_t, \theta_{i_t}))$. Let h_0 be the empty history. For each $t = 1, ..., n$, let $b(h_{t-1})$ be the probability distribution over the set $N \backslash \{0, i_1, ..., i_{t-1}\}$ that determines which of the remaining uninterviewed applicants will be interviewed in round t. Let $s(h_t)$ be the probability of stopping the interviewing in round t. The interviewing automatically stops in round n, so $s(h_n) = 1$. Let $a(h_t)$ be the probability distribution over the set $\{0, i_1, ..., i_t\}$ that determines which of the interviewed applicants is hired in the event that the procedure stopped in round t. Let $a_i(h_t)$ denote the probability that i is hired in that event, $i \in \{0, i_1, ..., i_t\}$.

We consider procedures that satisfy two properties. First, whenever the procedure stops, an interviewed applicant cannot be hired if another interviewed applicant is better. Formally,

$$\text{for each } h_t \text{ and each } i, j \in \{0, i_1, ..., i_t\}, \text{ if } \theta_i \prec \theta_j, \text{ then } a_i(h_t) = 0. \tag{A$_1$}$$

Second, whenever an applicant with the ideal job fit $\bar{\theta}$ is interviewed, the procedure stops and this applicant is hired, as there is nothing to gain by continuing the interviewing process. Formally,

$$\text{for each } h_t, \text{ if } \theta_{i_t} = \bar{\theta}, \text{ then } s(h_t) = 1 \text{ and } a_{i_t}(h_t) = 1. \tag{A$_2$}$$

As the probability of the ideal job fit $\bar{\theta}$ can be arbitrarily small, this assumption does not substantially constrain the applicability of our framework. Note that we implicitly assume that the job fit of applicant 0 is not ideal, that is, $\theta_0 \prec \bar{\theta}$, as otherwise the interviews would not even start.

Many real-life hiring procedures are included in our framework. Elements that can be incorporated in our model include interviewing one-by-one or in batches, as well as preferential treatment based on an applicant's identity i which incorporates observable information, such as gender and race. Different degrees of selectivity can be captured, from hiring the first acceptable applicant to interviewing many applicants before making a hiring decision.

Note that we assume *free recall*, in the sense that all of the applicants interviewed in earlier rounds remain available for hire. We also assume that the applicants never reject job offers. In Section 4.1 we show how the model and the results extend if we take possible unavailability of applicants and job offer rejections into account.

Importantly, we do not specify the employer-specific elements of the problem, namely, the employer's prior beliefs about the job fits and their cost of interviewing, as these should not play a role in understanding and evaluating fairness of a hiring procedure.

3. Fair Hiring

3.1. **Fairness.** A procedure is postulated to be fair if it satisfies two principles: *equal treatment of equals* and *invariance under reordering*. The former requires that, at all stages of the procedure, applicants with equivalent job fits have the same chance to be hired. The latter requires that the chance of any applicant to be hired should not change if any two interviewed applicants swap their positions in the order of the interviews. Note that the latter is equivalent to assuming that no reordering of interviewed applicants makes any difference. These principles must hold for every possible profile of the applicants' job fits.

To describe these principles formally, we use the following notation. Denote by $\boldsymbol{\Theta}$ the set of possible profiles of job fits of $n+1$ applicants, so $\boldsymbol{\Theta} = \{\theta_0\} \times \Theta^n$. Let $\boldsymbol{\theta} = (\theta_0, \theta_1, ..., \theta_n)$ be a profile of job fits, so $\boldsymbol{\theta} \in \boldsymbol{\Theta}$. Consider a procedure $\pi = (b, s, a)$.

Let $H^\pi(\boldsymbol{\theta})$ be the set of histories that can occur under π with a positive probability conditional on a given $\boldsymbol{\theta}$. For each round $t = 1, ..., n$ and each history h_{t-1} preceding that round, denote by $p_i^\pi(\boldsymbol{\theta}, h_{t-1})$ be the probability that applicant i will be hired (in round t or in a later round) conditional on the profile $\boldsymbol{\theta}$ of job fits, history h_{t-1}, and the event that the procedure reaches round t. So, $p_i^\pi(\boldsymbol{\theta}, h_{t-1})$ is evaluated at the start of round t.

We now state our first principle.

(P_1) *Equal Treatment of Equals.* For all profiles of job fits $\boldsymbol{\theta} \in \boldsymbol{\Theta}$, all histories $h_t \in H^\pi(\boldsymbol{\theta})$, and all applicants $i, j \in N$,

$$\text{if } \theta_i \sim \theta_j, \text{ then } p_i^\pi(\boldsymbol{\theta}, h_t) = p_j^\pi(\boldsymbol{\theta}, h_t).$$

Next we describe our second principle. Observe that the strategy b prescribing who to interview next after each history can be equivalently described as an ex ante strategy that, at the outset, randomly determines an order in which applicants are to be interviewed. Let $I = (i_1, ..., i_n)$ be an order over applicants in $N\backslash\{0\}$, where i_t is the applicant in position t in the order. Let $\mathcal{I}^\pi(\boldsymbol{\theta})$ be the set of all orders over $N\backslash\{0\}$ that are possible under π conditional on a profile of job fits $\boldsymbol{\theta}$, so $I = (i_1, ..., i_n) \in \mathcal{I}^\pi(\boldsymbol{\theta})$ if and only if $\prod_{t=1}^n b_{i_t}(h_{t-1}) > 0$.[2]

Next, given $\boldsymbol{\theta}$ and $I = (i_1, ..., i_n)$, let $r_t^\pi(\boldsymbol{\theta}, I)$ be the probability that round t is reached under procedure π conditional on that applicants are interviewed in order I, so

$$r_t^\pi(\boldsymbol{\theta}, I) = \prod_{t'=1}^{t-1}(1 - s(h_{t'})),$$

where $h_{t'} = ((i_1, \theta_{i_1}), ..., (i_{t'}, \theta_{i_{t'}}))$. Let $q_j^\pi(\boldsymbol{\theta}, I)$ be the probability that applicant j is hired under π conditional on that applicants are interviewed in order I, so

$$q_j^\pi(\boldsymbol{\theta}, I) = \sum_{t=t_j^*(I)}^{n} r_t^\pi(\boldsymbol{\theta}, I)s(h_t)a_j(h_t),$$

where $t_j^*(I)$ is the position of applicant j in the order I.

Given an order $I = (i_1, ..., i_n)$ and two positions t and t', denote by $I_{i_t \leftrightarrow i_{t'}}$ the order that is the same as I, except that applicants i_t and $i_{t'}$ swap their positions.

(P$_2$) *Invariance under Reordering.* For all profiles of job fits $\boldsymbol{\theta} \in \boldsymbol{\Theta}$, all orders $I \in \mathcal{I}^\pi(\boldsymbol{\theta})$, and all rounds t and t' such that $1 \leq t' < t \leq n$ and $r_t^\pi(\boldsymbol{\theta}, I) > 0$,

$$\text{if } I_{i_t \leftrightarrow i_{t'}} \in \mathcal{I}^\pi(\boldsymbol{\theta}), \text{ then } q_j^\pi(\boldsymbol{\theta}, I) = q_j^\pi(\boldsymbol{\theta}, I_{i_t \leftrightarrow i_{t'}}) \text{ for all } j \in N.$$

Definition 1. A hiring procedure is called *fair* if it satisfies (P$_1$) and (P$_2$).

3.2. **Categorization Procedures.** We show that a necessary and sufficient condition for a hiring procedure to be fair is that it is a categorization procedure.

A categorization procedure can be described as follows. Prior to starting interviews, divide possible job fits into two categories, called *strong* and *weak*. These categories are such that θ_0 is weak, $\bar{\theta}$ is strong, and no job fit that is categorized as strong is equivalent or worse than any job fit that is categorized as weak. Then, begin the interviewing. Each interviewing round begins by selecting with equal probability

[2]There is a simple way to resolve all randomness in the choice of an order at the outset without knowing $\boldsymbol{\theta}$. Let $(z_1, ..., z_{n-1})$ be i.i.d. draws from the uniform distribution on $[0, 1]$. These are realized at the outset, before the interviewing. For each history $h_t = ((i_1, \theta_{i_1}), ..., (i_t, \theta_{i_t}))$, we use the following notation. Let $b_i(h_t) = 0$ for each i who is already interviewed, i.e., $i \in \{0, i_1, ..., i_t\}$. For each $i = 0, 1, ..., n$, let $B_i(h_t) = \sum_{j=0}^{i} b_j(h_t)$, so $B_i(h_t)$ is the CDF of $b_i(h_t)$. The profile of realized values $(z_1, ..., z_{n-1})$ fully determines the choice who to interview after each history h_t by prescribing to interview i such that $B_{i-1}(h_t) < z_t \leq B_i(h_t)$.

one of the applicants who have not yet been interviewed. This applicant is then interviewed. If they fall into the strong category, then hire them on the spot. If instead they fall into the weak category and not all applicants have been interviewed, then proceed to the next round. If all applicants have been interviewed and none of them is strong, then hire one of them (which may be 0) subject to the following two conditions. Equivalent applicants are hired with the same probability. Hiring does not depend on the order in which applicants were interviewed.

We proceed with the formal definition.

Definition 2. A hiring procedure $\pi = (b, s, a)$ is called a *categorization procedure* if the following holds for all $\boldsymbol{\theta} \in \boldsymbol{\Theta}$. There exists a set of job fits $Y \subseteq \Theta$ that does not depend on $\boldsymbol{\theta}$ such that $\bar{\theta} \in Y$, $\theta_0 \notin Y$ and $\theta \not\precsim \theta'$ for all $\theta \in Y$ and all $\theta' \notin Y$. Initially, $b(h_0)$ is the uniform distribution over $N\backslash\{0\}$. Then, for each $h_t = ((i_1, \theta_1), ..., (i_t, \theta_t)) \in H^\pi(\boldsymbol{\theta})$, if $\theta_{i_t} \in Y$, then $s(h_t) = 1$ and $a_{i_t}(h_t) = 1$. If $\theta_{i_t} \notin Y$ and $t < n$, then $s(h_t) = 0$ and $b(h_t)$ is the uniform distribution over $N\backslash\{0, i_1, ..., i_t\}$. If $t = n$, then $a_i(h_n) = a_j(h_n)$ for all $i, j \in N$ such that $\theta_i \sim \theta_j$, and $a(h_n') = a(h_n)$ for all $h_n' \in H^\pi(\boldsymbol{\theta})$.

When job fits can be placed on a numerical scale (where a higher value means a better fit), so $\Theta \subset \mathbb{R}$, the categorization of job fits into strong and weak follows a *threshold strategy*. In this case, a categorization procedure takes a very simple form. Applicants are interviewed in a random order, where the first interviewed applicant whose value meets or exceeds a specified threshold $\bar{y}$ is hired. If all applicants have been interviewed and no value is equal to or greater than $\bar{y}$ is found, then the best applicant (up to tie breaking) is hired. Formally, set $Y = \{\theta \in \Theta : \theta \geq \bar{y}\}$.

At the opposite extreme, when no two job fits in $\Theta\backslash\{\bar{\theta}\}$ are comparable, then fairness imposes no discipline on how job fits are categorized, except that the ideal job fit is strong and the job fit of applicant 0 is weak. Moreover, fairness imposes no discipline on who is hired when all applicants have been interviewed.

Note that simultaneous search is also a particular categorization. Here, the employer commits to interview all applicants, unless an applicant with the ideal job fit $\bar{\theta}$ is found. The latter could have a vanishingly small probability. Formally, set $Y = \{\bar{\theta}\}$.

We now present the main result of this paper.

Theorem 1. *A hiring procedure is fair if and only if it is a categorization procedure.*

The proof is in Appendix A.

Remark 1. For simplicity of exposition, we have assumed that every interviewed applicant remains available until the procedure is over and always accepts the job offer. However, in reality, applicants sometimes become unavailable or may decline job offers. As shown in Section 4.1, this can be easily included in the model. To do this, we introduce an *unappointable* job fit — any applicant with this job fit will never be offered the job. When an applicant is no longer available or declines a job offer, the job fit of this applicant is replaced by the unappointable job fit. The hiring then proceeds as if this was the original job fit.

We discuss implications of Theorem 1 and then provide the intuition for the result.

3.3. **Implications.** We start by highlighting what fairness means for hiring practices.

1. Limited value of observable characteristics. Everyone in the pool has an equal chance to be interviewed. The order in which interviews take place and the decisions when to stop and who to hire may not depend on the applicants' observable characteristics. In particular, even if the observable characteristics contain some information about the job fit, they have to be ignored and cannot influence hiring until the job fit is revealed during the interview. For example, it is unfair to eliminate at the start all applicants with age above 40. It is unfair to interview the youngest applicants first. It is unfair to give a preferential treatment to the younger of two applicants who have the same job fit.

2. Commitment to hiring criteria. Before starting the interviewing process, the employer has to commit to a categorization of the applicants' job fits into who will be designated as strong and hired on the spot, and who will be designated as weak and only possibly hired when all applicants have been interviewed. In particular, the employer may not adapt their hiring criterion to what they have learned during interviews. For example, it is unfair to conduct several initial interviews to test the waters before settling on a hiring criterion. When an applicant with a strong job fit is interviewed, it is unfair to become optimistic about the market and postpone hiring this applicant to check if someone even stronger turns up. When several applicants with very poor job fit are interviewed, it is unfair to become disappointed about the market and stop the intervieweing with no hire when others are still waiting to be interviewed.

3. Transparency and accountability. As a categorization needs to be committed in advance, categorization procedures are transparent. This counteracts a major concern in hiring that decisions are not transparent, potentially concealing direct and indirect discrimination. In fact, we advocate that employers should preregister their criteria of how the attributes of applicants are categorized into "strong" and "weak". This makes it easier to detect, investigate, and prosecute cases of discrimination. The preregistration process can even ensure the random order of interviewing, by using a certified randomizing device to determine the order and mandating the employer to conduct interviews in that order.

4. Simplicity. A fair procedure is simple to implement. Where the job fits can be placed on a numerical scale, it is described by a single threshold that separates strong and weak applicants. This threshold is chosen and fixed ex ante. Changing this threshold during the interviewing process is not allowed. The recall of applicants that were interviewed in the past is also ruled out except when all applicants have been interviewed. This simplicity stands in contrast with optimal search without fairness constraints. In the case of independent values, the employer sets a threshold and searches for the first applicant with a job fit about this value, possibly decreasing the threshold over time and possibly recalling the applicants who were interviewed in the past (Weitzman, 1979). The case of correlated values is difficult and only solved in closed form under normal distributions with uncertain mean (DeGroot, 1968), and under Dirichlet priors (Rothschild, 1974).

Even without preregistration, the clear structure of categorization procedures leads to testable implications. For instance, if the interviewing is stopped before all applicants have been interviewed and no one is hired, or the hired applicant is not the one interviewed last, then this cannot be a fair procedure.

Despite the constraints imposed on hiring that are needed to ensure fairness, fair hiring procedures are richer than they might seem.

1. Flexible objective. The employer is free to choose which job fits belong to which category, as long as they respect the given partial order. The categorization can depend on the employer's objective function, their costs, and their prior beliefs. For example, the employer may wish to hire quickly, in which case they would label many job fits as strong. At the other extreme, the employer may wish to interview all applicants unless an exceptionally good applicant is found. In this case they can categorize only very few job fits as strong.

2. Affirmative action. Affirmative action can be implemented by formally including the attribute of belonging to the minority group into the job fit. In that case, applicants who belong to different minority groups will have different job fits and can be treated differently.

3.4. **Intuition and Counterexamples.** In this section we provide intuition behind the proof of Theorem 1 and support it with several examples of procedures that are unfair according to our principles.

It is easy to see that every categorization procedure is fair. To see why (P_1) holds, consider two applicants with equivalent job fits. As these job fits are equivalent, they belong to the same category. Suppose that both job fits are strong. As each of these two applicants has the same chance to be interviewed before the other, they have equal chances to be hired. Alternatively, suppose that the job fits of these applicants are weak. These applicants can only be hired after everyone has been interviewed, and equivalent job fits are treated symmetrically. To see why (P_2) holds, suppose first that there are no strong applicants in the pool. Then the procedure interviews everyone, and the order in which the interviews are conducted is not allowed to matter according to the definition of a categorization procedure. Alternatively, suppose that there is at least one strong applicant in the pool. Consider a history where a strong applicant was interviewed and hired. Then, swapping this applicant with any applicant interviewed earlier would make no difference, as the same strong applicant would still be hired, just earlier. Next, consider a history where a weak applicant was interviewed last. This is only possible if no strong applicant has been interviewed yet. Then, swapping the last interviewed applicant with anyone interviewed earlier would make no difference, as this does not change the position of the first strong candidate in the order.

Let us now sketch the argument for why every fair procedure is a categorization procedure. Consider a procedure that satisfies (P_1) and (P_2). Condition (P_1) states that any two applicants with equivalent job fits must be treated equally. Thus, the applicants' identities and observable characteristics may not influence any decisions in the procedure. In particular, this establishes the property that all must have equal chance to be interviewed. Next, define as strong all the job fits such that if an applicant with this job fit is interviewed in round 1, they would be hired on the spot. We now use (P_2) to argue that once a strong applicant is interviewed, the interviewing must be stopped and this applicant must be hired. Indeed, if this applicant is interviewed in some later round but hired with probability less then one, they would be better off by moving up the order of interviews to round 1 where they would be hired for sure. We next use (P_1) and (P_2) to show that a weak applicant cannot be hired, unless all applicants have been interviewed. Indeed, if a weak applicant, call him Joe, is interviewed in a some round after 1 and hired with positive probability, he would be worse off by moving up the order of interviews to round 1. To see why, note first that Joe cannot be hired in round 1 by the definition of being weak. Moreover, Joe cannot be hired in any later round $t < n$. This is because otherwise there may exist an uninterviewed applicant, call her Kate, whose job fit is equivalent to Joe's and who would be asymmetrically treated relative to Joe, as Joe is interviewed and hired with positive probability while Kate is not even interviewed yet. We thus establish that a fair procedure interviews applicants until finds the first strong applicant who is immediately hired. In the event that there are no strong applicants, all are interviewed. We then use (P_1) to show that equivalent applicants are hired with equal probabilities and (P_2) to show that the order in which the applicants are interviewed does not affect their choices to be hired.

To illustrate the role of conditions (P_1) or (P_2) in our theorem, we provide three examples of procedures that are not fair according to our principle. For this illustration, suppose that job fits can be placed on a numerical scale, so $\Theta \subset \mathbb{R}$.

Procedure A. Assume that the employer has an independent prior about the job fit of each applicant, and these priors are ordered in terms of the first order stochastic dominance. Then, within the model of Weitzman (1979), it is optimal to interview the applicants in this order. Weitzman's procedure violates condition (P_1), because two applicants with equivalent job fits have different positions in the interview order and are not given equal opportunity to demonstrate their fit and to get the job.

Procedure B. Consider the following procedure. Applicants are divided into several batches. The division can be based on the applicants' resumes, or it can be random. Then, all applicants in the first batch are interviewed. If none of their job fits is above a specified threshold, then all applicants in the next batch are interviewed, and so on. This procedure once again violates condition (P_1), because two applicants with equivalent job fits who happen to be sorted in different batches are not given equal opportunity to demonstrate their fit and to get the job.

Procedure C. Consider the optimal procedure for the secretary problem of Fox and Marnie (1960). Applicants are interviewed in a random order. The first k applicants are never hired. The interviews continue until an applicant with a job fit better than those among the first k is found. This applicant is hired on the spot. If no such applicant is found after all have been interviewed, then the outside option 0 is chosen. This procedure violates condition (P_2), because the applicants interviewed in the first k rounds are set to fail, simply because they are interviewed too early. Each of them could be better off if they switch their place with the one in position $k+1$.

4. Extensions

In this section we present two extensions.

4.1. **Unavailability of Applicants.** Suppose that interviewed applicants may become unavailable. There are two possible interpretations. First, an applicant may not be interested in the job any longer (for example, because they accepted a job elsewhere). Second, if offered the job, an applicant might reject the offer (for example, because they hold a better offer from another employer). The possibility that applicants become unavailable can be incorporated into our model as follows.

Introduce a job fit $\underline{\theta}$ that is worse than the outside option, so $\underline{\theta} \in \Theta$ and $\underline{\theta} \prec \theta_0$. Thus, by ($A_1$), applicants with this job fit will never be offered the job. We will call applicants with this job fit *unappointable*.

Suppose that in each round, each applicant (interviewed or not) may become unavailable. For each $j \in N$ and each $\theta_j \in \Theta$, the probability of the event that j becomes unavailable in any given round is $\lambda(\theta_j)$. It is independent of other applicants and identical across rounds. Assume $\lambda(\theta) \in [0, 1)$ for all $\theta \in \Theta$. If $\lambda(\theta) = 0$ for all θ, then we are back to the baseline model.

Notice that we can no longer interpret the outside option 0 as an existing employee or an applicant whose job fit is known at the outset. Indeed, if 0 was a person, they could still reject the job offer and walk away. So, we would still need an outside option to capture what would happen if everyone in N walked away. Thus, assume that 0 is the outside option, and θ_0 need not be interpreted as the job fit. Rather, let θ_0 be an additional element outside of the set of job fits Θ that is not equivalent to any job fit in Θ, and extend the domain of all relevant functions to $\Theta \cup \{\theta_0\}$. Finally, assume that 0 is always available, so $\lambda(\theta_0) = 0$.

Hiring proceeds as described in Section 2, but with a single alteration. In each round t, after an interview is completed but before a decision to stop or continue the interviewing is made, all applicants in $N \backslash \{0\}$ are checked if they become unavailable. If j becomes unavailable, which happens with probability $\lambda(\theta_j)$, then, from that moment onwards, the procedure treats j as unappointable. That is, j's job fit θ_j is replaced with $\underline{\theta}$ in the profile of job fits. After the availability check and replacements of the job fits as explained above, the procedure carries on as described in Section 2.

In this extended model, Theorem 1 holds with the same proof. To see why the proof does not change, observe that every history that could have occurred in the original model can also occur in the extended model. Thus, in the extended model, conditions (P_1) and (P_2) are stronger as they must hold for a larger set of histories. Consequently, the "only if" part of Theorem 1 holds. Moreover, it is easy to verify that the categorization procedures continue to satisfy (P_1) and (P_2) in the extended model, so the "if" part of Theorem 1 also holds.

4.2. **Shortlisting.** Our principles of fairness require the employer to give all applicants equal chance and, potentially, to keep interviewing them until everyone is interviewed. In reality, this might not be possible. Employers may operate under budget constraint and unable to interview more than a fixed number of applicants. In this case, if we were to insist on our fairness standards, then the only way to implement equal treatment under budget constraint is not to interview anyone at all.

To deal with the above concern, in this extension, we minimally relax our fairness principles. Specifically, instead of applying the principle of equal treatment of equals to all applicants, we apply it to those who have a positive probability of being interviewed under some profile of job fits. Formally, let I^π be the set of applicants who may be interviewed under the procedure π for some realized profile of job fits:

$$I^\pi = \{j \in N : \text{there exist } \boldsymbol{\theta} \in \Theta^n \text{ and } h_t \in H^\pi(\boldsymbol{\theta}) \text{ such that } j \in \{0, i_1, ..., i_t\}\}.$$

As this only applies to some profile of job fits, there may be other profiles where such an applicant in I^π is not interviewed at all. An applicant who is not in I_π is never interviewed, no matter of what profile of job fits is realized.

We now define *weak equal treatment of equals*, which is the same condition as in (P_1), except that it applies only to applicants in I^π.

(P_1') *Weak Equal Treatment of Equals.* For all profiles of job fits $\boldsymbol{\theta} \in \boldsymbol{\Theta}$, all histories $h_t \in H^\pi(\boldsymbol{\theta})$, and all applicants $i, j \in I^\pi$,

$$\text{if } \theta_i \sim \theta_j, \text{ then } p_i^\pi(\boldsymbol{\theta}, h_t) = p_j^\pi(\boldsymbol{\theta}, h_t).$$

A hiring procedure is *weakly fair* if it satisfies (P_1') and (P_2). Then, the necessary and sufficient condition for a hiring procedure to be weakly fair is that it is a *shortlist-and-categorization* procedure. Any such procedure is described as follows. Prior to starting interviews, make two commitments. First, choose a subset of applicants called the *shortlist*. Second, divide possible job fits into two categories, *strong* and *weak*. These categories are such that θ_0 is weak, $\bar{\theta}$ is strong, and no job fit that is categorized as strong is equivalent or worse than any job fit that is categorized as weak. Then, follow a categorization procedure as defined in Section 3.2, with the difference that only the applicants in the shortlist are selected for interview, with equal probability.

We sketch the proof of the above claim. It is easy to verify that the shortlist-and-categorization procedures satisfy (P_1') and (P_2), so the "if" part of the claim holds. To prove the "only if" part, fix a procedure π that satisfies (P_1') and (P_2). Recall the notation I^π, which denotes the set of applicants who may be interviewed under π for some profile of job fits. Let us call I^π the shortlist. Since the applicants outside of I^π are never interviewed, w.l.o.g. we can reduce the pool of applicants to I^π. We then apply Theorem 1 to establish that the only procedures that are fair to this restricted pool of applicants are the categorization procedures.

Although our fairness conditions do not restrict the employer's choice of who to shortlist and who to eliminate, the key features of shortlist-and-categorization procedures are commitment and, hence, transparency. The employer is required to commit to a shortlist ex ante, before the procedure starts. If a shortlist is reported to a regulator, then it can be scrutinized, and cases of discrimination can be detected, investigated, and prosecuted.

5. Conclusion

Fairness and equal opportunity are at the centerstage of organizing modern society. Fairness in job hiring has received particular attention (see the literature surveyed in the introduction on biases when evaluating job fit). Yet there has not been much formal debate on whether the current practices in hiring are fair from the procedural perspective. One obstacle is that an adequate concept of procedural fairness has not been formulated for hiring procedures. A difficulty in doing this is that hiring is not a static situation but a dynamic process. Subtle details, such as the order of interviews or the role of applicants' resumes, influence whether someone might be hired. This makes it easy to disguise an unfair treatment. Moreover, unfairness can inadvertently emerge even if the employer intends to be fair.

By embedding the concern for fairness in a model of sequential search, we are able to formulate procedural fairness principles and establish their consequences for hiring. In particular, our analysis shows that it is easy to qualify whether any given hiring process is fair. We recall three examples. Equal treatment of equals is violated if applicants are interviewed in a predetermined order as in the optimal procedure of Weitzman (1979). Hiring the first applicant that outperforms an initial set of applicants (as optimal in the secretary problem of Fox and Marnie, 1960) violates invariance under reordering. The common practice of first interviewing a small batch and then adding more applicants when the first interviews are not successful is similarly not fair as it also violates invariance under reordering.

Our fairness principles not only reveal whether any given procedure is fair, they also allow us to identify that a hiring procedure is fair if and only if it can be described as a categorization procedure. These procedures are simple and transparent, and thus they are easy to monitor. Monitoring is important, as a categorization procedure can still be misused. Ideally, a hiring procedure should be preregistered with a regulator, to help detect and challenge in court any misuse such as the one in the example above.

Appendix A. Proof of Theorem 1

It is straightforward to verify, as outlined in Section 3.4, that every categorization procedure satisfies (P_1) and (P_2). We now prove that every procedure π that satisfies (P_1) and (P_2) is a categorization procedure.

Let $\pi = (b, s, a)$ be a procedure that satisfies (P_1) and (P_2). To show that π satisfies Definition 2 for all $\boldsymbol{\theta} \in \boldsymbol{\Theta}$ and all $h_t \in H^\pi(\boldsymbol{\theta})$, it will be convenient to use the following notation instead. Let $\bar{H}^\pi$ be the set of all histories that are possible under π, so

$$\bar{H}^\pi = \bigcup\nolimits_{\boldsymbol{\theta} \in \boldsymbol{\Theta}} H^\pi(\boldsymbol{\theta}).$$

Given a history $h_t \in \bar{H}^\pi$, let $\boldsymbol{\theta}_t$ be the profile of job fits of the applicants who are interviewed under history h_t, so $\boldsymbol{\theta}_t = (\theta_0, \theta_{i_1}, ..., \theta_{i_t}) \in \{\theta_0\} \times \Theta^t$. Let $\boldsymbol{\theta}_{-t}$ be a profile of job fits of the applicants who are not yet interviewed under history h_t, so $\boldsymbol{\theta}_{-t} \in \Theta^{n-t}$. Note that $\boldsymbol{\theta}_t$ is pinned down by h_t, whereas $\boldsymbol{\theta}_{-t}$ is arbitrary. We will now prove that π satisfies Definition 2 for all $h_t \in \bar{H}^\pi$ and all $\boldsymbol{\theta}_{-t} \in \Theta^{n-t}$, where the profile of job fits $\boldsymbol{\theta}$ is implicitly defined as $\boldsymbol{\theta} = (\boldsymbol{\theta}_t, \boldsymbol{\theta}_{-t})$.

We first prove several lemmas. The first lemma shows that applicants must be interviewed in random order. That is, an interviewee is chosen equally likely among those applicants who have not yet been interviewed.

Lemma 1 (Random Order)**.** *For all $h_t \in \bar{H}^\pi$ with $t < n$,*

$$b_j(h_t) = \frac{1}{n-t} \quad \textit{for each } j \in N \backslash \{0, i_1, ..., i_t\}. \tag{1}$$

Proof. Let $h_t \in \bar{H}^\pi$ with $t < n$. First, suppose that $t = n - 1$, so only one uninterviewed applicant is left in N. Denoting this applicant by j, we have $b_j(h_t) = 1$, so (1) holds. Next, suppose that $t \leq n - 2$. We show that $b_{j'}(h_t) = b_{j''}(h_t)$ for any two uninterviewed applicants j' and j'', and, thus, (1) follows from the definition of $b(h_t)$. Since $b(h_t)$ is independent of $\boldsymbol{\theta}_{-t}$, let all uninterviewed applicants have the ideal job fit, so $\boldsymbol{\theta}_{-t} \in \Theta^{n-t}$ satisfies $\theta_i = \bar{\theta}$ for all $i \in N \backslash \{0, i_1, ..., i_t\}$. Then, by ($\mathrm{A}_2$), whoever is interviewed in round $t + 1$ will have the ideal job fit and must be hired on the spot. Thus, for any $j', j'' \in N \backslash \{0, i_1, ..., i_t\}$ we have $p^\pi(j'|\boldsymbol{\theta}, h_t) = b_{j'}(h_t)$ and $p^\pi(j''|\boldsymbol{\theta}, h_t) = b_{j''}(h_t)$. Applying ($\mathrm{P}_1$) to j' and j'' yields $p^\pi(j'|\boldsymbol{\theta}, h_t) = p^\pi(j''|\boldsymbol{\theta}, h_t)$. Thus, we obtain $b_{j'}(h_t) = b_{j''}(h_t)$. $\square$

The second lemma shows that the stopping decision after every history is deterministic, so the procedure never randomizes between stopping and continuing.

Lemma 2 (Deterministic Stopping)**.** *For all $h_t \in \bar{H}^\pi$ with $t \geq 1$, $s(h_t) \in \{0, 1\}$.*

In the proof of Lemma 2 and thereafter, we will use the following notation. Given a history h_t and an uninterviewed applicant $j \in N \backslash \{0, i_1, ..., i_t\}$, let $h_t \oplus (j, \theta_j)$ be the history that follows h_t by interviewing applicant j in round $t + 1$.

Proof. For $t = n$ we have $s(h_n) = 1$ by definition. Let $h_t \in \bar{H}^\pi$ with $t < n$. By contradiction, suppose that $s(h_t) \in (0, 1)$. Since $s(h_t)$ is independent of $\boldsymbol{\theta}_{-t}$, let an uninterviewed applicant $j \in N \backslash \{0, i_1, ..., i_t\}$ have the job fit $\bar{\theta}$, so $\boldsymbol{\theta}_{-t}$ is an arbitrary element of Θ^{n-t} such that $\theta_j = \bar{\theta}$. Since $s(h_t) < 1$, the history $h_t \oplus (j, \theta_j)$ occurs with positive probability. Thus, there exists an order $I \in \mathcal{I}^\pi(\boldsymbol{\theta})$ whose first $t+1$ applicants are $(i_1, ..., i_t, j)$. Since $\theta_j = \bar{\theta}$, by (A_2), after interviewing j, the procedure stops and hires j. We thus have

$$q_j^\pi(\boldsymbol{\theta}, I) = \prod\nolimits_{t'=1}^{t} (1 - s(h_{t'})) \leq 1 - s(h_t) < 1,$$

where the last inequality is by $s(h_t) > 0$. Now, consider the order $I_{j \leftrightarrow i_1}$ which is the same as I except j and i_1 are swapped. So, j is interviewed first. Again, since $\theta_j = \bar{\theta}$, by (A_2) we have

$$q_j^\pi(\boldsymbol{\theta}, I_{j \leftrightarrow i_1}) = 1.$$

But (P_2) requires $q_j^\pi(\boldsymbol{\theta}, I) = q_j^\pi(\boldsymbol{\theta}, I_{j \leftrightarrow i_1})$. We reached a contradiction with (P_2). □

The third lemma shows that when the procedure stops and chooses which applicant to hire, applicants with equivalent job fits are hired with equal probabilities.

Lemma 3 (Ex Post Equal Treatment)**.** *For all $h_t \in \bar{H}^\pi$ with $t \geq 1$ and $s(h_t) = 1$,*

$$a_i(h_t) = a_j(h_t) \quad \textit{for all } i, j \in \{0, i_1, ..., i_t\} \textit{ such that } \theta_i \sim \theta_j.$$

Proof. Consider $h_t \in H^\pi(\boldsymbol{\theta})$ with $t \geq 1$ such that $s(h_t) = 1$. Use label k for the applicant interviewed in round t under h_t, so $k = i_t$. Since $a(h_t)$ is independent of $\boldsymbol{\theta}_{-t}$, let all uninterviewed applicants have the ideal job fit, so $\boldsymbol{\theta}_{-t} \in \Theta^{n-t}$ satisfies $\theta_i = \bar{\theta}$ for all $i \in N \backslash \{0, i_1, ..., i_t\}$. Let us evaluate the probability that $j \in \{0, i_1, ..., i_{t-1}, k\}$ is hired, from the perspective of the beginning of round t (before k is interviewed). If any applicant other than k is interviewed in round t, then j cannot be hired, because the interviewed applicant has job fit $\bar{\theta}$ and, by (A$_2$), is hired on the spot. Applicant k is interviewed with probability $1/(n-t+1)$ by Lemma 1. If k is interviewed, then $s(h_t)=1$ by assumption, so the procedure stops and hires j with probability $a_j(h_t)$. Thus we have

$$p_j^\pi(\boldsymbol{\theta}, h_{t-1}) = \frac{1}{n-t+1} a_j(h_t). \tag{2}$$

For all $i, j \in \{0, i_1, ..., i_{t-1}, k\}$, by (P$_1$), $\theta_i \sim \theta_j$ implies $p_j^\pi(\boldsymbol{\theta}, h_{t-1}) = p_i^\pi(\boldsymbol{\theta}, h_{t-1})$. Consequently, by (2), we obtain that $\theta_i \sim \theta_j$ implies $a_i(h_t) = a_j(h_t)$. □

The fourth lemma shows that if the procedure stops before all applicants have been interviewed, then the last interviewed candidate is hired. So, the procedure does not "recall" earlier applicants, unless everyone is interviewed.

Lemma 4 (No Recall)**.** *For all $h_t \in \bar{H}^\pi$ with $1 \le t < n$, if $s(h_t) = 1$ then $a_{i_t}(h_t) = 1$.*

Proof. Consider $h_t \in \bar{H}^\pi$ with $1 \le t < n$. Denote by j the last interviewed applicant, so $j = i_t$. Suppose that $s(h_t) = 1$. By contradiction, suppose that someone other than j may be hired, so there exists $k \in \{0, i_1, ..., i_{t-1}\}$ such that $a_k(h_t) > 0$. Since $s(h_t)$ and $a(h_t)$ are independent of $\boldsymbol{\theta}_{-t}$, let an uninterviewed applicant $i \in N \backslash \{0, i_1, ..., i_{t-1}, j\}$ have a job fit equivalent to k's, so $\boldsymbol{\theta}_{-t}$ is an arbitrary element of Θ^{n-t} such that $\theta_i \sim \theta_k$. Let us compare the probabilities of hiring k and i evaluated at the beginning of round t after history h_{t-1}. By the time the procedure stops at or after round t, there are three possibilities. First, i is interviewed. Then, k and i must be hired with the same probability by Lemma 3. Second, i is not interviewed and k is not hired. Then neither k nor i are hired. Finally, i is not interviewed and k is hired. This event occurs with probability at least $s(h_t)a_k(h_t)/t$, which is positive by assumption. It follows that k is strictly more likely to be hired than i, i.e., $p_k^\pi(\boldsymbol{\theta}, h_{t-1}) > p_i^\pi(\boldsymbol{\theta}, h_{t-1})$. As $\theta_i \sim \theta_k$, we reached a contradiction to (P$_1$). □

We now prove that π is a categorization procedure. Let Y be the subset of job fits in Θ defined as follows. The procedure π stops in round 1 with certainty if and only if the job fit of the first interviewee is revealed to be in Y. Formally:

$$Y = \{\theta \in \Theta : s(j, \theta) = 1 \text{ for all } j \in N \backslash \{0\}\}.$$

The proof is divided into four steps.

The first step shows that whenever an applicant is interviewed and their job fit is in Y, the procedure stops with certainty.

Step 1. *For all $h_t \in \bar{H}^\pi$ with $t \geq 1$, if $\theta_{i_t} \in Y$ then $s(h_t) = 1$.*

Proof. Consider an applicant $j \in N$ with $\theta_j \in Y$ and a history $h_t \in \bar{H}^\pi$ where j is interviewed last, so $(i_t, \theta_{i_t}) = (j, \theta_j)$.

If $t = 1$, then $s(i_1, \theta_{i_1}) = 1$ by definition of Y.

Consider $t \geq 2$. By contradiction, suppose that $s(h_t) < 1$. So, by Lemma 2, $s(h_t) = 0$. Let I be an order in $\mathcal{I}^\pi(\boldsymbol{\theta})$ whose first t applicants are $(i_1, ..., i_{t-1}, j)$. Since $s(h_t)$ is independent of the job fits of the uninterviewed applicants, let $\boldsymbol{\theta}_{-t}$ be such that $\theta_i = \bar{\theta}$ for all $i \in N \backslash \{0, i_1, ..., i_t\}$. Since $s(h_t) = 0$, the procedure continues to round $t+1$. Since an applicant interviewed in round $t+1$ will have the ideal job fit $\bar{\theta}$, this applicant will be hired by (A_2). So j cannot be hired under order I. We thus have

$$q_j^\pi(\boldsymbol{\theta}, I) = 0.$$

Now, consider the order $I_{j \leftrightarrow i_1}$ which is the same as I except j and i_1 are swapped. So, j is interviewed first. Since $\theta_j \in Y$, by definition of Y we have

$$q_j^\pi(\boldsymbol{\theta}, I_{j \leftrightarrow i_1}) = 1.$$

But (P_2) requires $q_j^\pi(\boldsymbol{\theta}, I) = q_j^\pi(\boldsymbol{\theta}, I_{j \leftrightarrow i_1})$. We reached a contradiction with (P_2). □

The second step shows that whenever an applicant is interviewed and their job fit is not in Y, and not everyone has been interviewed yet, the procedure continues to the next round.

Step 2. *For all $h_t \in \bar{H}^\pi$ with $t < n$, if $\theta_{i_t} \notin Y$ then $s(h_t) = 0$.*

Proof. First, consider round 1. Fix a job fit $\theta \notin Y$. By definition of Y, there exists an applicant $j \in N \backslash \{0\}$ such that $s(j, \theta) < 1$. So, by Lemma 2, $s(j, \theta) = 0$. By contradiction, suppose that there exists $k \in N \backslash \{0, j\}$ such that $s(k, \theta) > 0$. So, by Lemma 2, $s(k, \theta) = 1$. Since $s(j, \theta)$ and $s(k, \theta)$ are independent of the job fits of applicants other than j and k, let $\boldsymbol{\theta}$ be such that $\theta_j = \theta_k = \theta$, and $\theta_i = \bar{\theta}$ for all $i \in N \backslash \{0, j, k\}$. By ($A_2$), whenever an applicant with job fit $\bar{\theta}$ is interviewed, this applicant is hired on the spot. Thus, k may be hired in two cases: if k is interviewed in round 1 (by Lemma 4) and if j is interviewed in round 1 and k is interviewed in round 2. In contrast, since $s(j,\theta)=0$, j may only be hired in round 2 after both k and j are interviewed. In summary, we have

$$p_k(\boldsymbol{\theta},h_0)=\frac{1}{n}+\frac{1}{n(n-1)}s((j,\theta),(k,\theta))a_k((j,\theta),(k,\theta)),$$
$$p_j(\boldsymbol{\theta},h_0)=\frac{1}{n(n-1)}s((j,\theta),(k,\theta))a_j((j,\theta),(k,\theta)).$$

Since $\theta_j=\theta_k=\theta$, we have $a_j((j,\theta),(k,\theta))=a_k((j,\theta),(k,\theta))$ by Lemma 3. But then $p_k(\boldsymbol{\theta},h_0)>p_j(\boldsymbol{\theta},h_0)$, which contradicts ($\mathrm{P}_1$).

Next, consider a round $t\in\{2,...,n-1\}$, an applicant $j\in N$ with $\theta_j\notin Y$, and a history $h_t\in\bar{H}^\pi$ where j is interviewed last, so $(i_t,\theta_{i_t})=(j,\theta_j)$. By contradiction, suppose that $s(h_t)>0$. So, by Lemma 2, $s(h_t)=1$. Let I be an order in $\mathcal{I}^\pi(\boldsymbol{\theta})$ whose first t applicants are $(i_1,...,i_{t-1},j)$. Since $s(h_t)$ is independent of the job fits of the uninterviewed applicants, let $\boldsymbol{\theta}_{-t}$ be such that $\theta_i=\bar{\theta}$ for all $i\in N\backslash\{0,i_1,...,i_t\}$. Since $s(\cdot)$ is deterministic by Lemma 2 and history h_t is possible by assumption, the procedure reaches round t with certainty conditional on order I. Moreover, since $s(h_t)=1$, by Lemma 4, j is hired with certainty conditional on order I, so

$$q_j^\pi(\boldsymbol{\theta},I)=1.$$

Now, consider the order $I_{j\leftrightarrow i_1}$ which is the same as I except j and i_1 are swapped. So, j is interviewed first. Since $\theta_j\notin Y$, we have $s(j,\theta_j)=0$, as proven above. Then, by Lemma 3, j cannot be hired unless all applicants are interviewed. But then, since there are uninterviewed applicants with the ideal job fit $\bar{\theta}$, by (A_2) we have

$$q_j^\pi(\boldsymbol{\theta},I_{j\leftrightarrow i_1})=0.$$

But (P_2) requires $q_j^\pi(\boldsymbol{\theta},I)=q_j^\pi(\boldsymbol{\theta},I_{j\leftrightarrow i_1})$. We reached a contradiction with (P_2). □

The third step shows that if all applicants' job fits are not in Y (so everyone is interviewed with certainty by Step 2), then which applicant is hired does not depend on the order of interviews.

Step 3. *If $\theta_i\notin Y$ for all $i\in N$, then $a(h_n)=a(h'_n)$ for all $h_n,h'_n\in H^\pi(\boldsymbol{\theta})$.*

Proof. Suppose that $\theta_i\notin Y$ for all $i\in N$. Consider any two distinct histories $h_n,\tilde{h}_n\in H^\pi(\boldsymbol{\theta})$. Suppose that, by contradiction, there exists $j\in N$ such that $a_j(h_n)>a_j(\tilde{h}_n)$. Since $\tilde{h}_n$ can be obtained from h_n by a finite sequence of swaps between pairs of applicants, there exist two histories h'_n and h''_n which differ by a single swap of applicants i' and i'' such that $a_j(h'_n) > a_j(h''_n)$. Let I be the order of interviews under h'_n, and let $I_{i' \leftrightarrow i''}$ be the order of interviews under h''_n. By Step 2, round n is reached with certainty conditional on both I and $I_{i' \leftrightarrow i''}$. Thus, we have $q^\pi_j(\boldsymbol{\theta}, I) = a_j(h'_n)$ and $q^\pi_j(\boldsymbol{\theta}, I_{i' \leftrightarrow i''}) = a_j(h''_n)$. Since $a_j(h'_n) > a_j(h''_n)$, we have $q^\pi_j(\boldsymbol{\theta}, I) > q^\pi_j(\boldsymbol{\theta}, I_{i' \leftrightarrow i''})$. We reached a contradiction with (P_2). □

The last step shows that the set Y has the properties specified in Definition 2.

Step 4. *Set Y satisfies $\bar{\theta} \in Y$, $\theta_0 \notin Y$, and if $\theta \in Y$ and $\theta' \notin Y$ then $\theta \not\preceq \theta'$.*

Proof. First, $\bar{\theta} \in Y$ by (A_2) and the definition of Y.

Second, to prove that $\theta_0 \notin Y$, suppose by contradiction that $\theta_0 \in Y$. That is, there exists an applicant j with job fit $\theta_j = \theta_0$ such that $s(j, \theta_j) = 1$. By Lemma 1, the history $h_1 = (j, \theta_j)$ has positive probability. By Lemma 4, $a_j(j, \theta_j) = 1$, so j is hired with certainty. But, since $\theta_j = \theta_0$, by Lemma 3, j and 0 must have equal chances to be hired. We reached a contradiction.

Third, to prove that if $\theta \in Y$ and $\theta' \notin Y$, then $\theta \not\preceq \theta'$, consider applicants j and k with $\theta_j \in Y$ and $\theta_k \notin Y$, so $s(k, \theta_k) = 0$ by Step 2. By contradiction, suppose that $\theta_j \preceq \theta_k$. By Lemma 1, history $h_2 = ((k, \theta_k), (j, \theta_j))$ occurs with positive probability. By Step 1, $s(h_2) = 1$, and by Lemma 4, $a_j(h_2) = 1$, so j is hired with certainty. But if $\theta_j \prec \theta_k$, then k must be hired by (A_1). Alternatively, if $\theta_j \sim \theta_k$, then j and k must have equal chances to be hired by Lemma 3. We reached a contradiction. □

In summary, by Steps 1–4, a procedure π that satisfies (P_1) and (P_2) is a categorization procedure (Definition 2). This completes the proof of Theorem 1.

References

Arrow, K. J., A. Sen, and K. Suzumura (2010): *Handbook of Social Choice and Welfare*, vol. 2, Elsevier.

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.

Bolton, G. E., J. Brandts, and A. Ockenfels (2005): "Fair Procedures: Evidence from Games Involving Lotteries," *Economic Journal*, 115, 1054–1076.

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.

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., J. D. Gaebler, H. Nilforoshan, R. Shroff, and S. Goel (2024): "The measure and mismeasure of fairness," *Journal of Machine Learning Research*, 24, 14730–14846.

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.

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

Kagel, J. H. and A. E. Roth, eds. (1995): *The Handbook of Experimental Economics*, Princeton University Press.

Kleinberg, J., S. Mullainathan, and M. Raghavan (2017): "Inherent trade-offs in the fair determination of risk scores," in *Proceedings of the 8th Conference on Innovations in Theoretical Computer Science*.

Lazear, E. P., K. L. Shaw, and C. T. Stanton (2018): "Who Gets Hired? The Importance of Competition Among Applicants," *Journal of Labor Economics*, 36, S133–S181.

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

Mitchell, S., E. Potash, S. Barocas, A. D'Amour, and K. Lum (2021): "Algorithmic Fairness: Choices, Assumptions, and Definitions," *Annual Review of Statistics and Its Application*, 8, 141–163.

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

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.

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.