Fair Hiring Procedures

Karl H. Schlag
University of Vienna

Andriy Zapechelnyuk
University of Edinburgh

2023

1 What is this paper about?

Fairness, nondiscrimination, and equal opportunity are major concerns in hiring

Not about: Behavioral and statistical biases that may arise when assessing candidates (e.g., survey by Bertrand and Duflo 2017).
[Becker 1957, Phelps 1972]

About: Procedural fairness

In this paper, we:
- model hiring as sequential search;
- define procedural fairness;
- characterize which procedures are fair;
- discuss practical implications of fair hiring;
- evaluate the loss in efficiency due to fairness concerns.

2 Formal Model

A pool of n candidates, $n \geq 2$.

Candidates are interviewed in a given order.

Each candidate is identified by:
- his position $i = 1, 2, ..., n$ in the order of interviews,
- a profile of observable characteristics x_i (CV)
- a profile of attributes θ_i (fit for the job, type)

x_i are given

θ_i are (possibly non-i.i.d., possibly correlated) r.v.s

Outside option $i = 0$ with known type θ_0

3 Hiring Procedure

Hiring proceeds in rounds $t = 0, 1, ..., n$

- In round 0, the searcher decides whether to walk away with the outside option, or to begin the search.
- In each round $i = 1, ..., n-1$, candidate i is interviewed, and then 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.

4 Hiring Strategy

A hiring strategy prescribes when to stop the search and whom to hire

A pair $s = (\sigma, \varphi)$.

Let $h_t = (\theta_1, ..., \theta_t)$ be a history of types of interviewed candidates up to round t

A pair $s = (\sigma, \varphi)$:
- $\sigma(h_t)$ the probability of stopping;
- $\varphi(h_t)$ the probability distribution over $\{0, 1, ..., t\}$.

Note that there is free recall.

5 Assumptions

1. Types θ_i have common support Θ.

2. There is a candidate's type that is impossible to beat. Whenever such a candidate is interviewed, the searcher stops the interviewing process and hires this candidate. Formally:

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*.

6 Fairness

When is the procedure unfair? Three examples.

Criterion of fairness: Observable characteristics or the order in which interviews are conducted should not play any role for those who have revealed their true type
- Relies on the notion of permutation (multilateral swap of identities)

No permutation of candidates can increase the ex-ante probability that any interviewed candidate is hired, or the the ex-ante probability that outside option is kept.

7 Fair Strategies

THEOREM 1: A hiring strategy is fair if and only if it is a partition strategy.

Partition strategy:

- Specifies a subset of the possible types of candidates. Candidates whose types are in this subset are called *strong*.
- Candidates are sequentially interviewed until either a strong candidate is found or all candidates have been interviewed.
- If a strong candidate is found, then interviews stop immediately, and this candidate is hired on the spot.
- If all are interviewed, then the choice of whom to hire must not depend on the order in which the candidates were interviewed, and candidates who have the same type must be hired with the same probability.

8 Intuition

It is easy to see that every partition strategy is fair.

Let us sketch the converse.

1. Partition types according to the choice in the 1st round.
2. If a candidate with a strong type is discovered in any round, the search must stop immediately, and this candidate must be hired.
3. In every round, the search either stops or continues for sure.
4. An order where some non-strong candidate can be hired before everyone is interviewed must not exist.
5. An order where the outside option is chosen before everyone is interviewed must not exist either.
6. Finally, if the search reaches the last round, candidates with equal types must be treated equally. Moreover, who is hired must not depend on the order of interviews.

9 Implications

Commitment to the hiring criterion. Before starting the interviewing process, the hiring committee has to agree on a categorization of the candidates' attributes.

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.

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.

10 Implications (cont.)

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.

11 Extension: Fair Hiring Using Thresholds

Assume that the searcher's preferences over candidates' types admit a utility representation.

$v(\theta)$ the searcher's added value from hiring a candidate with type θ.

Normalize $v(\theta_0) = 0$.

Additional condition:
(C) Only the candidates with the highest value can be hired, ties resolved equally.

COROLLARY 1: A hiring strategy is fair and satisfies condition (C) if and only if it is a threshold strategy.

Threshold strategy: Partition by a threshold

12 Cost of Fairness

How much potential value the searcher loses by restricting herself to fair strategies?

Searcher's utility:
- u_0 is the baseline utility that the searcher obtains if she decides to stop in round $t = 0$ and choose the outside option
- Each type θ has value $v(\theta)$
- Cost $c \geq 0$ of each interview
- Discounting of future payoffs: 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

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

13 Cost of Fairness (cont.)

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.

Let C_{u_0} be the set of pairs of the interview cost c and discount factor δ that satisfy the above constraint.

14 Cost of Fairness (cont.)

Searcher's uncertainty:
- Ω has two states of the world (e.g., "good market" and "bad market")
- 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 $\lambda_\omega(\cdot|x_i)$
- Prior β is 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.

Assume that candidates who are positioned earlier in the order are likely to have higher values:

$$F_\omega(\cdot|x_i) \succeq_{fosd} F_\omega(\cdot|x_j) \text{ whenever } i < j. \tag{2}$$

Let $\mathcal{B}_{fosd}$ be the set of priors that satisfy condition (2).

15 Cost of Fairness (cont.)

We now define and compare optimal payoffs.

Given a prior $\beta \in \mathcal{B}_{fosd}$, a cost parameter c, and a discount factor δ:
- $U^*(\beta, c, \delta)$ is the maximal expected payoff when choosing among all hiring strategies.
- $U^P(\beta, c, \delta)$ and $U^T(\beta, c, \delta)$ are the maximal expected payoffs when choosing among partition and threshold strategies, respectively.

The next theorem provides tight upper bounds on the relative cost of fairness.

Theorem 1

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

16 Conclusion

Fair hiring is simple to implement

Fair hiring: choose upfront who will be considered as strong and who will not. Hire the first strong candidate on the spot.

The concern for fairness impose a relatively small cost, bounded by the factor of 2.