Sequential Search with Flexible Information *

PAVEL ILINOV[†]

École Polytechnique Fédérale de Lausanne

Andrei Matveenko[‡]

dérale de Lausanne University of Mannheim

SALIL SHARMA[§] Princeton University ELIAS TSAKAS[¶] Maastricht University MARK VOORNEVELD

December 2024

Abstract

We consider a model of sequential search in which an agent (the employer) has to choose one alternative (a candidate) from a finite set. A key feature of our model is that the employer is not restricted to specific forms of information acquisition, i.e., she is free to endogenously choose any interview for each candidate that arrives. Our main characterization result shows that the employer's unique optimal strategy is to offer a gradually easier interview to later candidates. Further, we study whether the candidates are treated equally in terms of the probability of being hired. We show that although the discrimination created by the order of consideration depends on the cost structure, for large numbers of candidates the first one will be generically favored. Finally, we characterize a wide range of cases where the employer prefers to start by interviewing ex ante worse candidates.

JEL-CODES: D81, D83, D91, J71.

KEYWORDS: sequential search, rational inattention, discrete choice.

^{*}We are grateful to Nageeb Ali, Georgios Angelis, Yair Antler, Sarah Auster, Luigi Butera, Mark Dean, Francesc Dilmé, Pierre Fleckinger, Duarte Gonçalves, Pavel Kocourek, Eugen Kovac, Daniel Krähmer, Jimmy Martinez-Correa, Filip Matějka, Fabio Michelucci, José L. Moraga, Marek Pycia, David Ronayne, Rastislav Rehák, Max Riegel, Nicolas Schutz, Alexander Sebald, Rava da Silveira, Vasiliki Skreta, Yiman Sun, Karol Szwagrzak, Yiannis Vailakis, Adrien Vigier, Mark Whitmeyer and numerous participants at various conferences and seminars for helpful comments and valuable feedback. Funding by the German Research Foundation (DFG) through CRC TR 224 (project B03) is gratefully acknowledged.

[†]Transport and Mobility Laboratory, École Polytechnique Fédérale de Lausanne, Switzerland (pawa.ilinov@gmail.com)

[‡]Department of Economics, University of Mannheim, Germany (matveenko@uni-mannheim.de) [§]Department of Politics, Princeton University, USA (ss5924@princeton.edu)

[¶]Department of Microeconomics and Public Economics, Maastricht University, The Netherlands (e.tsakas@maastrichtuniversity.nl)

Department of Economics, Stockholm School of Economics, Sweden (mark.voorneveld@hhs.se)

1 Introduction

Since the seminal work of Stigler (1961), search models have had a prominent place both within the microeconomic and macroeconomic literature (e.g., see Anderson and Renault, 2018; Chade et al., 2017; Mortensen and Pissarides, 1999). Multiple applications can be studied within this framework, e.g., a consumer's search through different alternatives before purchasing one, an employer's search through different candidates before hiring one, an employee's search through different jobs before accepting one, etc. The common denominator in all of these situations is that the search involves some kind of friction, which makes the search problem non-trivial.

Within this literature, sequential search models are considered a benchmark setting for studying such problems (McCall, 1970; Weitzman, 1979). In this basic environment, an agent goes through a number of options, acquiring information about them sequentially before stopping at some point and selecting one. While this model describes a somewhat stylized search procedure, it is still parsimonious enough to accommodate most applications of interest and thus allows us to study the most interesting features of search problems. In this setting, frictions usually arise as the cost of information acquisition or as waiting costs.

In this paper, we focus on sequential search, with the important novel feature that the agent can flexibly differentiate the information she collects about each option. Our analysis is carried out in the specific setting of an employer who sequentially interviews candidates for a vacancy, without being exogenously restricted to a specific set questions or a specific test format. The only constraint is that she must immediately decide on whether to hire the candidate or not after the interview. Such constraint is common in thick labor markets where decisions are made quickly and there are plenty outside options both for the employer and the candidates, e.g., low-skill workers in retail, hospitality sector, food services, construction, sales, etc. In fact, in many of these examples, work is temporary or seasonal and the workers are unlicensed which makes it practically unlikely to delay the hiring decision. At the same time, despite the large supply or candidates, employers face a hard deadline to hire someone, and as a result there is an exogenously given finite number of interviewing slots.

The reason why we are interested specifically in the role of flexible information acquisition is twofold: first, flexibility in the interviewing process is quite common in practice, as employers often keep calibrating their questions throughout the interviewing process (Laker, 2023); second, by allowing full flexibility, we can study the employer's best case scenario, thus obtaining an upper bound on what the employer can achieve in terms of efficiency. The latter will allow us to address questions regarding the optimal order of interviewing candidates (Weitzman, 1979) or regarding fairness concerns and biases that stem from the interviewing process (Bertrand and Duflo, 2017; Neumark, 2018; Schlag and Zapechelnyuk, 2024).

Our formal model is a variant of Pandora's boxes (Weitzman, 1979) with a rationally inattentive employer (Sims, 2003). In the benchmark case, the type of each candidate is drawn independently from the same (Bernoulli) distribution. Of course, the actual types remain unobservable to the employer. The candidates arrive for interviews sequentially in a fixed order. A candidate's interview takes the form of a usual Bayesian signal, which is chosen by the employer and crucially may depend on the realizations of the earlier interviews. Aligned with the rational inattention literature, the cost of each interview is posterior-separable (Caplin et al., 2022; Denti, 2022): in this sense, our model covers as special cases all the commonly-used cost specifications, including entropic costs. Upon observing the realized signal of an interview, the employer may either reject the candidate and proceed to the next interview or hire him right away, similarly to the literature on committed Pandora's boxes (Beyhaghi and Cai, 2023, Section 2.2) and on the secretary's problem (Correa et al., 2024, and references therein), as well as in the more extensive search literature within labor economics (Mortensen and Pissarides, 1999). As we have already discussed, this assumption naturally fits the applications that we have in mind, and moreover will eventually allow us to obtain sharp theoretical predictions.

We start by asking a general question: what is the employer's optimal interviewing strategy? We characterize the optimal strategy by reducing the dynamic information acquisition problem into a static one (Theorem 1). This characterization uncovers a special feature of the employer's behavior: we show that candidates interviewed earlier face a "more difficult interview." That is, the optimal interviews are ranked with respect to a novel order of difficulty. Remarkably, this order turns out to have striking similarities with the likelihood ratio dominance order for lotteries (Shaked and Shanthikumar, 2007, and references therein), the information bias order (Gentzkow et al., 2014), and a toughness measure in the context of product testing (Gill and Sgroi, 2012).

The implications of the previous result are twofold. First, the employer has the luxury of overshooting for very high expected quality in early candidates, as there are still plenty of candidates to come. Second, by having a large expected quality of a failed early candidate, the employer makes the interview of this early candidate relatively cheap, thus balancing the risk that she undertakes (via overshooting) and the information acquisition cost that she has to incur.

Our paper continues with another question with interesting practical implications for recruiters, candidates, regulators, etc. We ask whether sequential search with flexible information acquisition is discriminative, in the sense of ex-ante identical candidates having non-equal chances of being hired (Bertrand and Duflo, 2017; Neumark, 2018). Opinions among practitioners on this subject wildly differ, typically based on psychological arguments, i.e., some hiring managers suggest that it is better to be interviewed first because of primacy bias. In contrast, others suggest that it is better to go last because of recency bias (Cotterell, 2019; Laker, 2023). Here, we approach this problem from a completely different angle. We highlight that the optimal interviewing strategy induces a tradeoff: early candidates have a higher probability of being interviewed in the first place, but at the same time, they have a lower probability of being hired conditional on being interviewed because of the more difficult interview they face. This raises the question: which effect ultimately dominates?

We demonstrate that the answer to the question depends on the specific form of the learning cost. Additionally, the answer is significantly influenced by the number of candidates considered by the manager. We categorize our analysis into two sections: the "small world" where two candidates are considered, and the "big world" where many candidates are considered. For the small world, we provide a characterization for a broad class of functions, which includes entropy, quadratic costs, and Shorrock's entropies. For the big world, we show that the first arriving candidate will generically get preferential treatment and has the highest probability of being hired. That is, summing up, the candidates have the same probability of being chosen in very rare cases. A similar conclusion is reached in a very different model by Schlag and Zapechelnyuk (2024), who show that only very specific interviewing protocols guarantee fair treatment of the candidates.

In the final main part of the paper, we relax the assumption that candidates are ex-ante identical. Within this more general framework, we first naturally ask whether our results about the employer's optimal interview carry, and second, we ask whether the employer prefers to start by interviewing better or worse candidates first. Once again, similarly to the question about discrimination, this is a question to which practitioners do not agree (Selby Jennings, 2023). Starting with the first of these two last questions, we show that our previous result about the difficulty of the interview still holds, i.e., earlier interviews are easier than later ones. Then, turning to the second question in a setting with two candidates, we show that under quadratic costs, the order does not matter. However, with entropic costs, for large priors, it is often better to start with the better candidate, whereas for small priors, it is often better to start with the worse candidate.

Finally, we consider a simple extension of our model. Namely, we discuss a naive policy that forbids discrimination in the interview process: we force the employer to choose the same interview for all candidates. We show that in this case, the optimal interview that is used for all candidates is easier than the interview for the first candidate in the unrestricted case but more difficult than the interview for the second-to-last candidate. The consequence is immediate: the probability of hiring the first candidate in the restricted case is higher than the same probability in the unrestricted case. In many cases, the outcome of the naive policy can be contrary to the desired result: if the first candidate is given preferential treatment without restrictions, under the restricted interview policy, that candidate is treated even better.

Our results are important for several strands of literature. First, our work should be primarily seen as part of the literature on Pandora's boxes (Weitzman, 1979), and in particular on the case with no-recall (Salop, 1973). The main difference to our setting is that in Pandora's boxes, the information acquisition boils down to an allor-nothing decision, i.e., in our language, the employer will either learn the type of candidate with certainty (perfectly informative interview) or will not learn anything at all (completely uninformative interview). In addition, the early papers did not allow a candidate to be hired without having been interviewed first, something which was later permitted by Doval (2018). For an excellent recent overview of this literature, we refer to Beyhaghi and Cai (2023).

Within this literature, particularly relevant to our work are the recent papers of Schlag and Zapechelnyuk (2024), where the problem of fair sequential interviews is studied, and the one of Ursu et al. (2020), which allows for flexible information acquisition but restricts attention to a specific information acquisition technology.

Related to this literature is the one on the secretary's problem (Correa et al.,

2024, and references therein). This literature can be seen as a stream of sequential search with unknown distribution of types. The feature shared with our work is that they also typically assume no recall.

Second, our work can be seen as part of a broad stream within the dynamic rational inattention literature that focuses on the timing of information acquisition (Steiner et al., 2017; Morris and Strack, 2019; Zhong, 2022; Hébert and Woodford, 2023). There is variation in the underlying assumptions that they impose, e.g., some allowing for flexible information acquisition, some assuming discounting, and some considering continuous time. At the same time, all these papers differ from ours in that they allow for information acquisition about the entire state space at any point in time. In our context, this would imply that the employer can potentially interview multiple candidates at any point in time and may also invite back candidates for follow-up interviews. While this assumption certainly makes sense for the applications that these authors have in mind, it seems less appealing in the job market setting that we have in mind as our main application.

Third, our paper is incidentally related to the literature on ordering Bayesian signals, by introducing our difficulty order, which is equivalent to the toughness order that Gill and Sgroi (2012) use in a different context. A similar order has also been used by Gentzkow et al. (2014) and Charness et al. (2021) in an attempt to model information biases. All of these orders share striking similarities with the likelihood ratio dominance order (Shaked and Shanthikumar, 2007).

Finally, there is a large literature on ordered consumer search, studying dynamic information acquisition about different consumption choices before an eventual purchase decision is made (for an overview, see Armstrong, 2017). The main difference to our work is that the focus of this literature is on the role of different asymmetries across choices, e.g., which product does the consumer inspect first when the products differ in price or in inspection costs? Similarly to the literature on Pandora's Boxes, most of the work on ordered consumer search imposes strict exogenous assumptions on the information acquisition technology, which at the outset seems quite natural in the context of the corresponding applications. To the best of our knowledge, there are two exceptions: the paper of Jain and Whitmeyer (2021) where the effect of flexible information acquisition on market outcomes is studied, and the one of Georgiadis-Harris (2024) where —among else— dynamic information acquisition with a fixed deadline is studied, and it is found that close to the deadline not much is learned.

2 Model

We study a (female) employer who considers an ordered set of a priori identical (male) candidates $I = \{1, ..., T\}$. Each candidate $i \in I$ is associated with a type

$$\theta_i \in \Theta := \{ \text{Good}, \text{Bad} \} = \{ G, B \},\$$

which is independently drawn from the same distribution that assigns probability $\mu \in (0, 1)$ to the good type G^{1} .

 $^{^{1}}$ The binary type assumption is not essential for our analysis. Our results hold in a more general setting, more specifically, when the utility of the employer only depends on the beliefs

The employer must choose one candidate, and there is no outside option. Before making a decision, she may acquire information about the candidates' types. Information acquisition is sequential, following the candidates' order. That is, at stage i, the employer selects a Blackwell experiment $\sigma_i : \Theta \to \Delta(S_i)$ for candidate i. Upon observing a signal realization $s \in S_i$, she forms a posterior distribution on Θ , that we identify with her posterior belief about the state G

$$p_i^s := \frac{\mu \sigma_i(s|G)}{\mu \sigma_i(s|G) + (1-\mu)\sigma_i(s|B)}.$$

It follows from the work of Kamenica and Gentzkow (2011) that each experiment is identified by a mean-preserving distribution over posterior beliefs, i.e., by some $\pi_i \in \Delta([0, 1])$ such that $\mathbb{E}_{\pi_i}(p) = \mu$. A fully uninformative signal is one that puts probability 1 to the prior μ .

After updating to belief p_i^s about candidate *i*, the employer either hires *i* or proceeds to interview the next candidate i + 1. We assume no recall, i.e., if a candidate is not hired right after an interview, he is no longer available to the employer. There can be several rationales for such an assumption: For instance, it can be driven by psychological factors of the rejected agent (e.g., pride, frustration, etc.) or by conditions on the labor market (e.g., there is excess labor demand and the rejected candidate is hired immediately by another firm), or because of institutional rules (e.g., HR rules dictate that every candidate is only interviewed once). Thus, the employer chooses an action

$$a_i \in A_i := \{0, 1\}$$

following the realization of an interview for candidate i, where action 0 corresponds to not hiring a candidate and 1 to hiring. Hiring a good candidate brings utility 1, and hiring a bad candidate 0.

Formally, a non-terminal history at round $i \in \{1, 2, ..., T\}$ is identified by the set of realized posteriors for all candidates $j \in \{1, ..., i-1\}$, i.e.,

$$\mathcal{H}_i := [0,1]^{i-1}.$$

The employer's action at every $h \in \mathcal{H}_i$ consists of a signal that leads to a distribution of posteriors π_i^h and a mapping α_i^h : $\operatorname{supp}(\pi_i^h) \to A_i$. Whenever $p \in \operatorname{supp}(\pi_i^h)$ is realized, and $\alpha_i^h(p) = 1$ is chosen, a terminal history is reached, and candidate *i* is hired. In case round *T* is reached, the last candidate will be hired regardless of the realization of the respective signal. A typical strategy of the employer is henceforth denoted by $(\boldsymbol{\pi}, \boldsymbol{\alpha})$.

Information acquisition is costly. In line with the rational inattention literature, we consider posterior separable costs (Caplin et al., 2022; Denti, 2022): test π_i costs

$$C(\pi_i) = \lambda \bigg(\mathbb{E}_{\pi_i}[c(p)] \bigg), \tag{1}$$

about posterior means as is typically assumed in the literature of the information design or costly information acquisition, for example, see Arieli et al. (2023) and Mensch and Malik (2023) for the recent references. We hold the binary type assumption mainly for the simplicity of the exposition.

where $\lambda \in \mathbb{R}_{++}$ is the marginal cost of information and $c : [0, 1] \to \mathbb{R}$ is continuous, strictly convex and smooth on the interior of the unit interval.

Throughout the paper, we impose the following assumption:

BOUNDARY CONDITION :
$$\lim_{p \to 0^+} c'(p) + \frac{1}{\lambda} < c'(\mu) < \lim_{p \to 1^-} c'(p) - \frac{1}{\lambda}.$$
 (2)

Intuitively, this assumption puts a lower bound on the marginal cost at the boundaries, i.e., learning the true state with certainty is expensive.

The two standard special cases of costs are the quadratic and the entropic costs.

Example 1. We say that the cost function is *quadratic* whenever

$$c(p) = (p - \mu)^2$$

Example 2. We say that the cost function is *entropic* whenever

$$c(p) = \mu \log \mu - (1 - \mu) \log(1 - \mu) - (-p \log p - (1 - p) \log(1 - p)).$$

Recall that $-p \log p - (1-p) \log(1-p)$ is known as Shannon entropy.

If the employer chooses the action (π_i^h, α_i^h) at history $h \in \mathcal{H}_i$, the expected payoff that she wants to maximize is equal to

$$\mathbb{E}_{\pi_i^h} \Big[\alpha_i^h(p) p + (1 - \alpha_i^h(p)) V_i - \lambda c(p) \Big],$$

where V_i denotes her maximum net expected payoff if she continues and interviews candidate i + 1. Note that V_i depends only on the number of remaining candidates, as the types of the different candidates are drawn independently from the same probability distribution, and there is no recall possibility. Hence, without loss of generality, we can restrict attention to \mathcal{H}_i -measurable strategies, i.e., to strategies such that $(\pi_i^h, \alpha_i^h) = (\pi_i, \alpha_i)$ for all $h \in \mathcal{H}_i$. This means that the employer's expected payoff is simplified to

$$\mathbb{E}_{\pi_i} \Big[\alpha_i(p) p + (1 - \alpha_i(p)) V_i - \lambda c(p) \Big].$$
(3)

Definition 1. The employer's dynamic optimization problem is

$$\max_{(\pi_i,\alpha_i)} \mathbb{E}_{\pi_i} \Big[\alpha_i(p) p + (1 - \alpha_i(p)) V_i - \lambda c(p) \Big],$$
(4)
subject to

$$V_{i} = \max_{(\pi_{i+1},\alpha_{i+1})} \mathbb{E}_{\pi_{i+1}} \Big[\alpha_{i+1}(p)p + (1 - \alpha_{i+1}(p))V_{i+1} - \lambda c(p) \Big],$$
(5)

$$V_T = 0. (6)$$

Constraint (5) ensures dynamic consistency; that is, the DM behaves optimally in every history. Constraint (6) captures the intuition that if the DM has rejected all candidates, she ends up with zero payoff.

 \triangleleft

 \triangleleft

3 Optimal interviewing strategy

The optimal strategy (π, α) in problem (4) is the collection of the optimal actions (π_i, α_i) for all $i \in I$. Note that the interview design problems at some stages i, j differ only by their continuation values V_i and V_j , correspondingly. Therefore, in the dynamic problem, the employer behaves as if she solves a collection of static problems with different continuation values. These continuation values are determined from the future behavior of the employer, and the value is *exogenous* at stage i. Thus, at stage i, a continuation value V_i serves the role of an outside option to the employer. Therefore, we conclude that at each stage i, the employer solves a static problem with an exogenous outside option. A static problem with an exogenous outside option is a building block for the dynamic problem, and we discuss a static problem in great detail in this section.²

Additionally, we observe that at stage *i* given the realized value *p*, the employer simply selects candidate *i* if $p \ge V_i$ and continues search otherwise, and therefore in our previously-stated optimization problem, we can replace $\alpha_i(p)p + (1 - \alpha_i(p))V_i$ with $\max\{p, V_i\}$. Thus, the optimization problem at stage *i* boils down to the following (static) optimization problem with parameter $V := V_i$.

Definition 2. The static optimization problem (or the problem with exogenous outside option) is

$$\max_{\pi} \mathbb{E}_{\pi} \Big[\underbrace{\max\{p, V\} - \lambda c(p)}_{\phi(p, V, \lambda)} \Big].$$
(7)

We discuss the properties of the static problem using two technical lemmas. The first lemma characterizes the solution exploiting the convexity and differentiability of the function c(p).

Lemma 1. The following statements hold:

- 1. The solution to problem (7) exists, and it is unique and interior.
- 2. There exist two thresholds $V_L, V_H \in (0,1)$ with $V_L < \mu < V_H$, such that the optimal test π_V satisfies:

$$V \leq V_L \qquad \Longrightarrow \qquad \operatorname{supp}(\pi_V) = \{\mu\}, \\ V_L < V < V_H \qquad \Longrightarrow \qquad \operatorname{supp}(\pi_V) = \{p_V^L, p_V^H\}, \\ V \geq V_H \qquad \Longrightarrow \qquad \operatorname{supp}(\pi_V) = \{\mu\}.$$

3. The optimal hiring decision is given by the following:

| $V \leq V_L$ | | | \implies | $\alpha(\mu) = 1,$ |
|-----------------|-----|-------------|------------|--------------------|
| $V_L < V < V_H$ | and | $p = p_V^H$ | \implies | $\alpha(p) = 1,$ |
| $V_L < V < V_H$ | and | $p = p_V^L$ | \implies | $\alpha(p) = 0,$ |
| $V \ge V_H$ | | | \implies | $\alpha(\mu) = 0.$ |

 $^{^{2}}$ Such a problem is a variant of the problem of a rationally inattentive agent with an exogenous outside option, see, e.g., Matějka and McKay (2015), Caplin and Dean (2013) for the reference.

The previous lemma is illustrated in Figure 1. The idea is that $\phi(p, V, \lambda)$ consists of two strictly concave parts, with a kink at V. This induces the two posteriors p_V^L and p_V^H , and as the prior lies between these two, the employer will acquire an informative signal that distributes its probability to these two posteriors; otherwise, she will pick the completely uninformative signal. These posteriors are obtained, for example, using the concavification technique as in Caplin and Dean (2013).³



Figure 1: The employer's static payoff function and its concavification when $V \in (V_L, V_H)$.

The key observation is that both p_V^L and p_V^H are, in general, continuously increasing in V (see Lemma B4 in the Appendix for a formal proof). Moreover, we have

$$\lim_{V \to 0^+} p_V^L = \lim_{V \to 0^+} p_V^H = 0 \text{ and } \lim_{V \to 1^-} p_V^L = \lim_{V \to 1^-} p_V^H = 1.$$

Hence, for sufficiently large V, the whole interval $[p_V^L, p_V^H]$ will lie to the right of μ . Likewise, for sufficiently small V, the interval will lie to the left of μ . Thus, we can define the two thresholds:

$$V_H := \min\{V \in [0,1] : p_V^L \ge \mu\}, V_L := \max\{V \in [0,1] : p_V^H \le \mu\}.$$

Intuitively, if the value of the outside option is very high $(V > V_H)$, the DM does not acquire any information and simply chooses the outside option. Similarly, when the value of the outside option is very low $(V < V_L)$, the DM takes the current option without extra information. Note that by the prior $\mu \in (0, 1)$ being full-support, we obtain both $V_L \in (0, 1)$ and $V_H \in (0, 1)$.

In the dynamic problem of the employer, the outside option at stage i equals the value of the problem at stage i + 1. In turn, this value equals the maximal attained value in a static problem (7) for a particular value of outside option V. To

³For recent use of concavification to the related rationally inattentive problems, see, e.g., Jain and Whitmeyer (2021), Kim et al. (2022).

study the maximal attained value in the static problem, we define a value function $g:[0,1] \rightarrow [0,1]$ such that

$$g(V) = \max_{\pi} \mathbb{E}_{\pi} \Big[\phi(p, V, \lambda) \Big]$$

Function g(V) is clearly non-decreasing by construction, as $\phi(p, V, \lambda)$ is weakly increasing in V for every p. We note that function g(V) is linear on $[0, V_L]$ and $[V_H, 1]$, namely, we have $g(V) = \mu$ if $V \leq V_L$ and g(V) = V if $V \geq V_H$, as in either of these two regions the employer does not incur any costs for acquiring information, and makes a hiring decision straight away. The following lemma completes the analysis for the entire unit interval.

Lemma 2. The function g is increasing, convex, and differentiable on [0,1]. Moreover, it is strictly increasing on the interval $[V_L, 1]$.

We will now proceed to characterize the solution to the dynamic problem. To do so, we first define a specific sequence of static problems by means of a sequence of outside options, namely, for each $i \in \{1, \ldots, T\}$, we have

$$V_i := g(V_{i+1}),\tag{8}$$

with $V_T = 0$.

Lemma 3. For the sequence $(V_i)_{i=1}^T$ defined in (8), the following statements hold:

- 1. The continuation value V_i is strictly decreasing in i.
- 2. For every $i \in \{1, ..., T-1\}$, we have $\mu \leq V_i < V_H$.

We illustrate Lemma 3 graphically (Figure 2). First of all, by combining convexity and differentiability of g with the fact that g(V) = V for all $V \ge V_H$, it follows that g(V) > V for all $V < V_H$. Hence, V_i keeps shrinking as the employer moves to later candidates. This is not surprising, as there are fewer candidates left to interview. The fact that for all candidates except the last one, the outside option always lies in the learning region (V_L, V_H) holds similarly.

Theorem 1. The solution to the dynamic problem of Definition 1 is as follows:

- 1. At every round $i \in \{1, ..., T-1\}$, the employer draws the signal π_{V_i} which is optimal in the static problem (according to Lemma 1), with the outside option V_i that we defined in (8). Moreover, we have:
 - (a) If $p_{V_i}^H$ is realized, the search stops and candidate *i* is hired.
 - (b) If $p_{V_i}^L$ is realized, the search continues to candidate i + 1.
- 2. At round T, the employer does not acquire information and hires the candidate right away.



Figure 2: The sequence of outside options with three candidates.

The employer's optimal interview design in the dynamic problem is very straightforward: She continues the search until she receives a high signal about the quality of a candidate. If only the low signals have been realized during the first T - 1 interviews, she simply chooses the last candidate with the fully uninformative signal.

In the next section it will be easier to interpret optimal interview at stage i as a binary test. The employer offers personalized test to each candidate i. If the candidate passes, he is hired; if the candidate fails, he is discarded. We state our main result about the characterization of the optimal tests in the next section.

4 Difficulty of interviews

We introduce the following partial order to compare the different tests that the employer optimally chooses for different candidates.

Definition 3. Let π_i and π_j be two binary tests, with $p_i^L < p_i^H$ and $p_j^L < p_j^H$ being the posteriors beliefs in the respective supports. We say that π_i is more difficult than π_j if

$$p_i^L > p_j^L$$
 and $p_i^H > p_j^H$.

The condition above has a simple interpretation. The two candidates are ex ante identical from the point of view of the employer, and they are offered one test each such that, whenever there is a tie, i is deemed better than j, i.e., in particular,

(a) if both of them pass their respective tests, the expected quality of i is higher than the expected quality of j, and

(b) if both of them fail their respective tests, the expected quality of i is higher than the expected quality of j.

The concept of test difficulty can also be demonstrated by means of a simple example. Consider two following interview tasks, the first being "Formulate and prove the Hahn-Banach theorem" and another one "Compute the derivative of quadratic function" The first task is clearly more difficult, and if the candidate succeeds, the posterior belief about the candidate becomes relatively high, while if he fails, the belief does not go down by much. For the second task, which is clearly easier, the belief dynamics are the opposite.

Let us provide some further justification for calling π_i more difficult than π_j . Recall that the two tests π_i and π_j are respectively characterized by the underlying experiments σ_i and σ_j (see Section 2).

Definition 4. We say that experiment σ_i likelihood-ratio dominates experiment σ_j whenever, for every signal realization $s \in \{H, L\}$,

$$\frac{\sigma_i(s|G)}{\sigma_i(s|B)} > \frac{\sigma_j(s|G)}{\sigma_j(s|B)}.$$

The underlying idea is as follows: conditional on every test result, the relative evidence for the good type is stronger under i's interview than under j's interview. A similar relation has been used in the literature on biased information sources (Gentzkow et al., 2014; Charness et al., 2021) and in the literature on product testing (Gill and Sgroi, 2012). All these orders bear striking similarities with the likelihood ratio dominance, which is often used in the literature to compare lotteries (Shaked and Shanthikumar, 2007).

Proposition 1. The following are equivalent:

- (i) Test π_i is more difficult than test π_j .
- (ii) Experiment σ_i likelihood-ratio dominates experiment σ_j .
- (iii) For every prior μ , the passing probability under σ_i is lower than the passing probability under σ_i , i.e.,

$$\underbrace{\mu\sigma_i(H|G) + (1-\mu)\sigma_i(H|B)}_{passing \ probability \ under \ \sigma_i} < \underbrace{\mu\sigma_j(H|G) + (1-\mu)\sigma_j(H|B)}_{passing \ probability \ under \ \sigma_i}.$$

(iv) The probability of "false-positive" outcome $\sigma_i(H|B)$ is lower than $\sigma_j(H|B)$, and the probability of "false-negative" outcome $\sigma_i(L|G)$ is higher than $\sigma_i(L|G)$.

From the previous result, it follows directly that our notion of more difficult interview does not depend on the prior. This is a desirable property, satisfied by other well-known orders over the set of Bayesian experiments. The idea is that difficulty is a property of the test alone, defined independently of the candidate who takes the test. Furthermore, if the two candidates share the same prior, i's interview is more difficult than j's interview if and only if the probability of passing the test is lower for i than it is for j.

Let us now state our main result, which says that, in the optimal interviewing strategy, the interviews decrease in difficulty as the employer proceeds to later candidates.

Theorem 2. In the optimal strategy from Theorem 1, the following hold:

- 1. Difficulty of interviews is decreasing with respect to the order of being interviewed, i.e., for all $i \in \{1, ..., T-2\}$, test π_{V_i} is more difficult than test $\pi_{V_{i+1}}$.
- 2. As the number of candidates grows large, we obtain:

$$\lim_{T \to \infty} p_{V_1}^L = \mu \text{ and } \lim_{T \to \infty} p_{V_1}^H = V_H.$$

Hence, the following hold:

- (a) The probability of the first candidate being hired converges to 0.
- (b) The probability of hiring a good candidate is bounded away from 1.

Figure 3 shows an example of an optimal learning strategy. On the vertical axis, we have the number of remaining candidates besides the one currently interviewed. So, for instance, if there are ten candidates in total, we depict the optimal interviews for the first nine, recalling that the last one will be hired anyway without an interview.



Figure 3: Optimal interviews as a function of the number of remaining candidates.

Decreasing the high posterior realizations $p_{V_i}^H$ is intuitive. If a posterior $p_{V_i}^H$ is realized on the interview *i*, the employer stops the search and chooses candidate *i*. Thus, it makes sense that in order for the employer to stop the search early,

she needs to be sufficiently certain that the candidate she hires is good, as she is foregoing the chance to interview many other potential candidates.

Decreasing the low posterior realizations $p_{V_i}^L$ is less intuitive. By using such a strategy, the employer optimally procrastinates: instead of acquiring the most information during the first interviews, she wants to spread expected information acquisition towards all interviews. Intuitively, during the first interviews, she offers difficult tests to the applicants because she has some applicants left. The employer wants to bear the risk and try to "catch a big fish" at the beginning and do not spend too much on information to be more sure that the rejected candidate is of bad type. The fewer candidates are left, the safer the strategy used by the employer is.

Additionally, we describe the dynamics of the optimal interviews in terms of the statistical errors that the employer makes. Combining results from Theorem 2 and Proposition 1, we obtain that, in the first stages, the employer bears the risks, offers the hardest tests, and tolerates the false negatives, whereas, in the later stages, she plays safer, decreasing the probability of false negatives and increasing the probability of false positives.

Corollary 1. The probability of false positive outcome $\sigma_i(H|B)$ is decreasing in *i*, and the probability of false negative outcome $\sigma_i(L|G)$ is increasing in *i*.

It is remarkable that even with an arbitrarily large number of candidates, the employer will not be certain that a good candidate will be hired in the end. This is because even the first candidate's test (which will be very difficult due to the vast number of candidates that are still to follow) will not fully rule out the possibility of a false positive. The reason is that the marginal cost of information close to the boundary grows arbitrarily large, i.e., it becomes too expensive to try and split hairs at the top end of test results.

5 Discrimination

From Theorem 2, we know that the optimal interviews are more difficult for earlier candidates: *conditional* on being interviewed, later candidates get an easier test than earlier ones. However, from Theorem 1, we know that the probability that a candidate is interviewed at all is lower than that of his predecessors since it depends on those candidates failing their interviews. The aggregate effect is unclear. In this section, we ask which effect dominates, that is, what can we say about the *unconditional* probability of a candidate being hired?

In our analysis, we distinguish two cases depending on the number of available candidates. We argue that in the case when a few candidates are available, the pattern of discrimination is complex and non-trivially dependent on the shape of the cost function and the parameters of the model. However, when the number of candidates is large, the pattern of discrimination starts to emerge, and the model delivers a unique prediction.

5.1 Few candidates

In the section, we focus on the analysis when T = 2. In order to make the problem tractable, we make certain assumptions on the cost function. We start with the technical assumption of symmetry.

Definition 5. c is said to be symmetric about a point $z \in (0,1)$ if for any pair $p, r \in [0,1]$

$$|z - p| = |z - r| \implies c(p) = c(r).$$

We denote the set of symmetric functions by C_s

Note here that the point of symmetry, z, may or may not depend on the prior belief. Symmetry is a mild assumption that is satisfied, for example, for quadratic or entropic costs. See also Hébert and Woodford (2021), who derive symmetric cost using an axiomatic approach.

Next we put an additional structure on the shape of the cost. We require our cost to belong to one of the three families.

Definition 6. *c* belongs to either of the three following families:

$$\mathcal{C}_1 = \left\{ c \in \mathcal{C}_s \mid c' \text{ is concave on } (0, z) \right\},\\ \mathcal{C}_2 = \left\{ c \in \mathcal{C}_s \mid c' \text{ is linear on } (0, z) \right\},\\ \mathcal{C}_3 = \left\{ c \in \mathcal{C}_s \mid c' \text{ is convex on } (0, z) \right\}.$$

The symmetry defines how the cost function c behaves on the interval (z, 1). If $c \in C_2$, then c is simply a quadratic cost function. The shape of the derivative c'(p) is crucial because it imposes a structure on the costs associated with obtaining high p^H and low p^L posteriors. In the optimum, the employer balances the informativeness of both low and high posteriors, ensuring that the marginal costs for acquiring p^H and p^L are aligned. For instance, if c'(p) is concave for small beliefs, it means that improving information about the good state is more expensive than improving it about the bad state. This is due to the symmetry that makes c'(p) convex for high beliefs.

We say that candidate 1 is *favored* if the unconditional probability of hiring him exceeds 0.5. The next proposition identifies the behavior of the employer depending on which family the cost function belongs to.

Proposition 2. *1.* If $c \in C_1$, then

- 1) If $\mu > z$, then the first candidate is favored.
- 2) If $\mu = z$, then no candidate is favored.
- 3) If $\mu < z$, then the second candidate is favored.
- 2. If $c \in C_2$, then no candidate is favored.
- 3. If $c \in C_3$, then
 - 1) If $\mu > z$, then the second candidate is favored.

- 2) If $\mu = z$, then no candidate is favored.
- 3) If $\mu < z$, then the first candidate is favored.

Intuitively, the first candidate is favored if the low posterior belief is more informative about the bad state than the high posterior about the good state; if p^L is further away from μ than p^H . Family C_s puts structure on such informativeness trade-off. As an immediate corollary from Proposition 2, we get results for quadratic and entropic costs because quadratic costs belong to C_2 family and entropic costs belong to C_1 family with z = 0.5.

Corollary 2. If c is quadratic, no candidate is favored. If c is entropic, the first candidate is favored if $\mu > 0.5$, the second candidate is favored if $\mu < 0.5$, and no candidate is favored if $\mu = 0.5$.

Bartoš et al. (2016) find that endogenous attention leads to discrimination when candidates have different expected productivity. They show that the behavior of the employer depends on whether the market is *lemon-dropping* with high priors or *cherry-picking* with low priors. Proposition 2 suggests that discrimination may present even if the candidates are ex-ante the same, but the choice problem is sequential.

There is an ongoing debate in the literature on rational inattention, which learning cost function most adequately represents the choice of the individual. We show that the observed behavior can depend on non-trivial properties of the cost, but our results also provide a testable hypothesis for the cost of learning. Our result in Proposition 2 can be parametrized with the help of Shorrock's entropy, which was recently used in experiments (Dean and Neligh, 2023) and in theoretical work (Miao and Xing, 2024)). Given a posterior belief p for our binary setup, Shorrock's entropy is

$$c(p) = -\frac{1}{(\alpha - 1)(\alpha - 2)} \Big[1 - \Big(p^{2-\alpha} + (1-p)^{2-\alpha} \Big) \Big],$$

where $\alpha \in \mathbb{R}$ is a parameter. Shorrock's cost belongs to the Uniformly Posterior Separable costs family. Additionally, note that when $\alpha \to 1$, Shorrock's entropy is approaching the negative entropy, and when $\alpha = 0$, the costs are quadratic.

Shorrock's entropy is symmetric around z = 0.5, and the sign of c'''(p) on (0, z) is uniquely determined by the value of α . The following corollary reformulates Proposition 2 for Shorrock's entropy.

Corollary 3. 1. If $\alpha \in (-\infty, -1) \cup (0, \infty)$, then

- 1) If $\mu > 0.5$, then the first candidate is favored.
- 2) If $\mu = 0.5$, then no candidate is favored.
- 3) If $\mu < 0.5$, then the second candidate is favored.
- 2. If $\alpha = 0$, then no candidate is favored.
- 3. If $\alpha \in (-1, 0)$, then
 - 1) If $\mu > 0.5$, then the second candidate is favored.

- 2) If $\mu = 0.5$, then no candidate is favored.
- 3) If $\mu < 0.5$, then the first candidate is favored.

Our Corollary 3 produces a testable hypothesis. The hypothesis can be tested as follows. The parameter α can be estimated from the earlier choices as in Dean and Neligh (2023). After that, the subject faces the sequential task that we describe in the current paper. Depending on the estimated value of α , we have a unique prediction about the discrimination that can arise. Depending on the results, it is an empirical question of whether Shorrock's entropy adequately describes the behavior of the subject.

Adding one candidate substantively complicates the analysis. The complexity of the discrimination pattern increases because of the endogenous continuation value. For instance, as shown in Figure 4 below, the probabilities of the different candidates being drawn for prior $\mu = 0.15$ and different values of the parameter λ , we see that the results are a very mixed bag.



Figure 4: Non-monotonicity with three candidates and entropic costs. q denotes unconditional probability of being hired.

These results may explain why there is no consensus as to whether it is beneficial for a candidate to be interviewed early or late. In particular, the folk views among practitioners are split, with some arguing that it is better to be interviewed early on and some arguing that it is better to be interviewed late. The typical explanation that is used to support either of these arguments is that employers are biased, suffering from primacy and recency bias, respectively. However, our previous analysis suggests that all the reasons why there is such a plurality of views among practitioners can very well be attributed to the sensitivity of the discrimination to unobservable parameters, but it can nonetheless still be explained through the lens of a rational model.

5.2 Many candidates

If we switch our focus to cases with a large number of candidates, we observe that a particular pattern starts to emerge as the number of candidates increases, irrespective of unobservable parameters. Namely, the first candidate starts emerging as the one who is consistently favored. As Proposition 2 suggests, the behavior of c'(p) is important. First, we show that if c'(p) is linear, then the first candidates are favored.

Denote by q_{iT} the unconditional probability of candidate *i* being hired from a pool of *T* candidates. Whenever *T* is obvious from the context, with slight abuse of notation, we will omit it and simply write q_i .

Proposition 3. Let $c \in C_2$. For every prior $\mu \in (0, 1)$, for any fixed T > 2, the unconditional probability is strictly decreasing in the order of being interviewed, i.e.,

$$q_1 > q_2 > \cdots > q_{T-1} = q_T.$$

Recall that earlier candidates face interviews that are harder to pass, but they have a higher probability of being interviewed. If c''(p) is constant, the second effect will always dominate: it is always better to be among the earlier candidates. For general costs, the analysis is much more complicated. However, if the number of candidates is large, the first tests should not vary much, and therefore, it should also be better to be first. Moreover, later candidates should have a lower probability of being interviewed. Overall, intuitively, it should also be better to be first with a large number of candidates. The next theorem confirms the intuition.

Theorem 3. For every prior $\mu \in (0,1)$, there is some $T_0 > 2$ such that for every $T > T_0$ the first candidate is favored, i.e.,

$$q_1 > \max\{q_2, \ldots, q_T\}.$$

We identify several implications of Theorem 3. First, the results provide the normative tool for a candidate: in the ex-ante identical environment without private information and many competitors, it is always better to be interviewed first. Second, the results are crucial for policymakers concerned about fairness as they reveal the direction of discrimination. Third, our Theorem 3 gives a unique prediction about the behavior of the individual modeled by the posterior-separable cost of information. Denti (2022) argues that the assumption of posterior separability may not hold in certain decision problems. Given the unique prediction, it is an empirical question whether the assumption is always appropriate for our problem.

It is now safe to generically conclude that in models with flexible information acquisition, there will generally be discrimination. The latter highlights that the usual tradeoff (between efficiency and equity) arises in a rather surprising form in this setting, i.e., the employer's first best (which is the overall most efficient outcome) will only guarantee equality of opportunity for the candidates under a very restricted set of parameters. A result in the same direction was recently obtained by Schlag and Zapechelnyuk (2024), who showed that only a very specific set of interviewing protocols can guarantee fair treatment of the candidates. This is a feature that may make a designer or policy maker uncomfortable, as equal opportunity is often a desired property for a hiring procedure (Bertrand and Duflo, 2017; Neumark, 2018).

6 Different productivities

Throughout the text, we assumed that all candidates were ex ante identical. Although that may seem a substantial simplifying assumption, it allows us to investigate the effect of the sequential structure on the employer's incentives to acquire information, without needing to worry about confounding effects due to differential prior information. In this section, we will relax this assumption, thus allowing the manager to have different priors about the different candidates, e.g., due to different education, work experience, recommendations from past employers, etc.

In our analysis, we compare the information acquisition strategies for the candidates with different priors. In general, with posterior separable costs, the marginal costs of information may depend on prior beliefs. As we are mostly interested in the intertemporal trade-offs, it is natural to exclude cases when different prior beliefs may affect the optimal test through the direct cost effect. Consequently, we consider uniformly posterior separable costs in this section: we assume that function c(p) is independent of the prior belief.⁴

The first question that naturally arises then is whether our earlier results on the employer's optimal interview will still carry. It is not difficult to verify that the basic structure remains the same, i.e., the outside option V_i will still be decreasing across candidates, in the sense that removing an early candidate can only be (weakly) detrimental for the employer. The sequence of outside options $(V_i)_{i=1}^T$ is still obtained by Lemma 3, with the only difference that the function g is now prior dependent and therefore becomes

$$g_i(V) = \max_{\pi} \mathbb{E}_{\pi} \Big[\phi(p, V, \lambda) \Big],$$

where π is chosen from the distributions with mean μ_i instead of μ .

As a result, we can simplify the employer's optimization problem by first dropping some candidates. There are two types of candidates that we can drop:

- 1. The candidates that are preceded by someone who is hired without an interview, i.e., every j = 2, ..., T such that $\mu_i \ge p_{V_i}^H$ for some i < j.
- 2. The candidates that are skipped without an interview, i.e., every *i* such that $\mu_i \leq p_{V_i}^L$.

These candidates do not play any role in the optimization problem, as it is known ex ante that they will never be optimally hired, either because they follow someone who is much better than them, or because they are followed by someone who is much better than them, respectively.

The remaining candidates can be seen as the ones that have been shortlisted for an interview. Notice that the way a shortlist is obtained is dependent on the order we plan to conduct the interviews, i.e., a different order may very well lead to a different shortlist. Once a shortlist has been obtained, the problem looks very similar to the one we previously solved, and the main conclusions of Theorem 2 still hold.

⁴Quadratic costs can be equivalently defined as a member of uniformly posterior separable cost, for example, as Shorrocks entropy with $\alpha = 0$.

Proposition 4. Suppose that $\tilde{I} = \{1, \ldots, \tilde{T}\}$ is the ordered set of shortlisted candidates with $\tilde{T} \geq 2$. Then, the employer's optimal strategy satisfies the following:

- 1. Difficulty is decreasing with the order of being interviewed, i.e., for every $i = 1, \ldots, \tilde{T} 2$, it is the case that $p_{V_i}^H > p_{V_{i+1}}^H$ and $p_{V_i}^L > p_{V_{i+1}}^L$.
- 2. As the number of shortlisted candidates grows large, we obtain:
 - (a) The probability of the first candidate being hired converges to 0.
 - (b) The probability of hiring a good candidate is bounded away from 1.

However, the really interesting question within this framework is whether the employer should prefer to interview the best or the worst candidate first. This question is discussed among practitioners (Selby Jennings, 2023), but the issue remains still unsettled.

Given the complexity of the question, we will focus on the simplest possible context, i.e., a setting with only two candidates, and the employer's cost function coming from one of the common specifications. Let μ_1 and μ_2 be the priors of the two candidates respectively. Moreover, denote by U_1 and U_2 the employer's indirect net utility when candidate 1 is interviewed first and when candidate 2 is interviewed first, respectively. Obviously, if $\mu_1 = \mu_2$, the two utilities are identical, and the order does not matter.

Theorem 4. Suppose that the are only two candidates, i.e., we have $I = \{1, 2\}$. Then, the following hold:

- (a) QUADRATIC COST: The employer does not have a preference on the order of interviewing the candidates, i.e., $U_1 = U_2$ for any $\mu_1, \mu_2 \in M$.
- (b) ENTROPIC COST: For any large prior there is an even larger prior with which the employer would rather start; likewise, for any small prior there is a larger prior with which the employer would rather follow, i.e., there exists some $\mu_0 \in$ (0,1) such that:
 - 1) For every $\mu_1 > \mu_0$, there exists some $\mu_2 > \mu_1$ such that $U_1 < U_2$.
 - 2) For every $\mu_1 < \mu_0$, there exists some $\mu_2 > \mu_1$ such that $U_1 > U_2$.

Fix the order of candidates. Theorem 4 leverages the idea that if switching the order changes the employer's strategy from non-learning to active learning, then the initial order can not be optimal. In the case of quadratic costs, the optimal posteriors are linear in the outside option, and this effect is not present. The effect arises for more complex relationships between the optimal posteriors and an outside option. To get analytical results, it is convenient to have closed-form expressions for posterior beliefs, and therefore, we formulate our results for the entropic costs.

Theorem 4 suggests that the employer's preferences regarding the order of interviewing the candidates is highly sensitive with respect to the prior beliefs, meaning that unless we exogenously impose very strong assumptions on the primitive parameters of the model, it will not be possible to make strong predictions about the employer's optimal order of interviewing candidates. This conclusion reinforces the view that general results about the optimal interviewing order can only be obtained under strong structural assumptions Doval (2018). This is in contrast to the impression that many people had regarding the generality of the corresponding result of (Weitzman, 1979).

7 Extensions

7.1 Restricted interview design

At the optimum, the employer in our model fully leverages the flexibility of the interview design. She constructs different interviews depending on the serial position of a candidate. Utilizing the dynamic structure of the problem, the employer offers more difficult tests to the candidates who are arriving early. Such a difference in the treatment can create discrimination in the hiring outcomes, as discussed in Sections 5 and 6, and may be seen as normatively unfair. In this section, we instead consider the scenario in which the employer is exogenously restricted to using interviews with the same difficulty for all candidates, except for the last one.⁵

We will maintain some features of our main model for comparison purposes. In particular, we will still consider common prior and only binary interviews. Moreover, passing an interview will always lead to a candidate being hired, whereas failing the interview will always lead to a candidate being rejected.

We formulate the restricted manager's problem as a restricted version of problem (1).

Definition 7. The restricted dynamic employer's problem is to find (π, α) such that

$$(\pi, \alpha) \in \left(\arg \max_{\substack{(\pi_{\sup p \mid \pi \mid \leq 2, \alpha)}}} \mathbb{E}_{\pi} \left[\alpha(p_{i})p_{i} + (1 - \alpha(p_{i}))V_{i} - \lambda c(p_{i}) \right] \right)$$

subject to
$$V_{i} = \mathbb{E}_{\pi} \left[\alpha(p_{i+1})p_{i+1} + (1 - \alpha(p_{i+1}))V_{i+1} - \lambda c(p_{i+1}) \right],$$

$$V_{T-1} = \mu,$$

$$V_{T} = 0.$$

We emphasize two differences between the restricted and the unrestricted problems. First, in the restricted problem, the employer chooses a posterior distribution only once. Therefore, we omit index i. Additionally, there is a restriction on the cardinality of the support of the posterior distribution.⁶

It is convenient to analyze the restricted problem (9) explicitly using the unconditional choice probabilities. The restricted problem is a special case of a static

⁵To not put our employer at risk of not hiring anyone, we assume that the last arriving candidate can be hired without an interview.

⁶If the support of posterior distribution is not bounded, it is generally unclear whether Lemma 1 holds in the restricted setting. For example, anticipating the lack of choice in the following periods, the DM may leverage her flexibility in the first period and design a test with more than two realizations. This strategy may allow the optimal dynamic behavior to be squeezed into a single test. Such a more general problem is out of this project's scope and is a natural direction for future research.

problem (2) with a corresponding outside option. In Appendix A, we reformulate the static problem (2) as an equivalent problem, in which maximization is over the unconditional choice probabilities only, and show that there is a one-to-one mapping between unconditional choice probability and the optimal posteriors. Analogous analysis can be found in a recent paper by Fosgerau et al. (2023) and is similar to the expressing discrete static rational inattention problem as the log-sum as in Matějka and McKay (2015), Caplin et al. (2019).

The next proposition is our main result for this Section; it compares the solutions of the restricted and unrestricted problems.

Proposition 5. Let q^* and q_{1T} be the optimal probability of hiring the first arriving candidate in the restricted and unrestricted problems with T candidates, respectively. When T > 2, the inequality $q_{12} > q^* > q_{1T}$ holds.

Proposition 5 identifies that in the restricted problem, the employer chooses a test with intermediate difficulty: the test is easier than the hardest one, which is offered to the first candidate, and is more difficult than the easiest one, which is offered to the last interviewed candidate. Recall that the employer engages in more risky behavior in the unrestricted setting in the earlier stages. She chooses a test with a low probability of success because she can mitigate the failure in the future stages. Restriction on the interview design forbids such mitigation. The results in Proposition 5 are intuitive and naturally fit into the risk interpretation: the manager bears less risk in the earlier stages.

It is immediate from Proposition 5 that the restriction on interviews increases and decreases the probabilities of hiring the first and the last interviewed candidates, respectively. How the restriction influences the chances of other candidates being hired is generally unclear. Consider the natural goal of the authority: make probabilities of being hired not depend on the serial number of a candidate and, therefore, be equal to each other. A combination of the results from Proposition 5 with the analysis in Section 5 suggests that the policy introduced in this Section may *increase* discrimination that appears in the unrestricted problem. For example, it happens if the sequence q_{iT} decreases in *i*. The ideal policy must take into account the intertemporal incentives of the manager and not naively put a simple restriction on the feasible set of strategies.

7.2 Infinitely many candidates

The alternative way to think about our problem is within the context of sequential search models à la McCall (1970). The key modeling difference between this approach and ours is that they typically assume an infinite stream of candidates also without recall— and they solve for the stationary optimal search strategy (see Chade et al., 2017, for an overview of this literature). To make the employer's problem well-defined, we weight her future payoffs with the discount factor $\beta \in (0, 1)$.

Definition 8. The employer's dynamic optimization problem with an infinite amount

of candidates is

$$\max_{(\pi_{\infty},\alpha_{\infty})} \mathbb{E}_{\pi_{\infty}} \Big[\alpha_{\infty}(p)p + (1 - \alpha_{\infty}(p))\beta V_{\infty} - \lambda c(p) \Big],$$
(9)
subject to

$$V_{\infty} = \max_{(\pi_{\infty}, \alpha_{\infty})} \mathbb{E}_{\pi_{\infty}} \Big[\alpha_{\infty}(p) p + (1 - \alpha_{\infty}(p)) \beta V_{\infty} - \lambda c(p) \Big].$$
(10)

The solution to problem (9) always exists and is unique for high enough β . Indeed, using the value function, we can express constraint (10) as $V_{\infty} = g(\beta V_{\infty})$. From Lemma 2 value function g(V) is strictly increasing, and, additionally inequalities, $g(\beta V_H) < V_H$ and $g(\beta \mu) > \mu$ hold. Therefore, equation $V_{\infty} = g(\beta V_{\infty})$ always has a unique solution.

The employer has a stationary optimal rule in problem (9) because the solution to the static problem (7) is unique. Therefore, the employer offers the same test to each candidate at the optimum. The solution is strikingly different from the solution of the main problem (4) when a very large number of candidates are available. Additionally, due to discounting, the employer chooses a test with intermediate difficulty when an infinite number of candidates are available: the optimal test is easier than the hardest one, which is offered to the first candidate, and is more difficult than the easiest one, which is offered to the last interviewed candidate. Intuition from Section 7.1 applies, rejection becomes costly, and, therefore, the employer engages in the less risky behavior.

With an infinite amount of candidates, the part of the Theorem 2 still applies; the employer will not be certain that a good candidate will be hired. Additionally, it is clear that the candidates who are interviewed early will be favored.

8 Conclusion

As documented in the economic literature, see, e.g., Bertheau et al. (2023), hiring is difficult for firms, and one of the reasons is that the firms face time constraints while hiring candidates. This means that firms do not learn the potential workers' productivities perfectly (since it will take too long time) but instead acquire noisy information about those. In this paper, we model the process of sequential search with costly but flexible learning in each stage.

The hiring firm observes several candidates who arrive sequentially and can design interviews for each candidate individually. We show that the optimal learning strategy has a simple feature – the later the candidate appears (the higher the serial number she has), the easier questions she will be facing. That is, the optimal interviews are decreasing in their difficulty in time. However, it does not mean that the workers should try to be interviewed in the end since the probability of being hired as a function of time of arrival is not necessarily increasing.

Our paper is the first step in studying sequential search with flexible and endogenous information acquisition. Therefore, many research questions are left for the future. For instance, we study only the situation when the distribution of productivities is known to the employer. The problem of studying a similar problem with extra layer of learning about the workers' productivities is interesting and intriguing.

Another suggestion for future research is to consider a model similar to ours but with an opportunity for recall. We suspect that the decreasing difficulty property will remain present in this class of problems.

A Reformulation of the static problem

Instead of relying on the concavification technique to solve the static problem (2), we will introduce the unconditional probability q of the interview having a good outcome explicitly as a choice variable. This is in contrast, in the concavification technique, the optimal unconditional probability was derived as a function of the two posterior beliefs that satisfy Bayesian consistency.

For every $V \in [0, 1]$, define the function

$$U(q, V) := \max_{x,y} \Big[q\phi(x, V) + (1 - q)\phi(y, V) \Big],$$

subject to
$$\mu \le x \le 1,$$

$$0 \le y \le \mu,$$

$$qx + (1 - q)y = \mu.$$

Note that for notation simplicity, we have omitted λ from ϕ . Then, we can reformulate the static optimization problem as follows:

Lemma A1. For every $V \in [0, 1]$, the optimization problem

$$\max_{q \in [0,1]} U(q, V) \tag{A.1}$$

has a unique solution, henceforth denoted by q_V . Moreover, we have

$$g(V) = U(q_V, V). \tag{A.2}$$

Additionally, problem (A.1) is concave in q, with unique solution, and the interior solution of (A.1) is decreasing in V.

Proof. We first analyze the case in which inequality $p^H > \mu > p^L$ holds. That is, when $V \in (V_L, V_H)$. Using the fact that support of the optimal posterior distribution has no more than two points we rewrite the static problem (2) as

$$\max_{\substack{(q,p^H,p^L)\in[0,1]^3}} \{q(p^H - \lambda c(p^H)) + (1-q)(V - \lambda c(p^L))\}$$

s.t. $qp^H + (1-q)p^L = \mu,$
 $p^H \ge p^L.$ (A.3)

We denote the objective in the above problem as $\tilde{U}(q, p^H, p^L)$. Using the identity from the multivariable calculus we can write $\max_{q,p^H,p^L} \tilde{U}(q, p^H, p^L) = \max_q \max_{p^H,p^L} \tilde{U}(q, p^H, p^L)$. Therefore, problem (A.3) can always be solved sequentially finding optimal p^L, p^H given q and then optimize over q. To show the equivalence between problems, we need to show first that for given q, there is only one pair of optimal (p^L, p^H) and second that there is a unique optimal q.

To show that there exists a unique pair of optimal (p^L, p^H) given q we observe that the interior solution to the static problem (2) should satisfy necessary optimality conditions in problem (A.3). In particular, for given q optimal (p^L, p^H) should satisfy the system

$$\begin{cases} -\lambda c'(p^L) = 1 - \lambda c'(p^H) \\ qp^H + (1-q)p^L = \mu. \end{cases}$$

We show that given $q \in (0,1)$ the system has a unique solution (p^L, p^H) . We rewrite the first equation as $p^L = (c')^{-1}(c'(p^H) - \frac{1}{\lambda})$. This expression defines a function $p^L(p^H)$. Indeed, $c'(p^H)$ is increasing in p^H , therefore mapping $p^L(p^H)$ is also increasing and the mapping defines unique p^L for any p^H .

We rewrite the second equation as $q = (\mu - p^L)/(p^H - p^L)$. Simple algebra shows that the derivative of the right-hand side with respect to p^H is positive; therefore, the right-hand side is increasing in p^H . Therefore, for any given q, there exists unique p^H and, thus, for any given q, there exists unique pair (p^L, p^H) that solves the system above. Further, we write $p^L(q), p^H(q)$ to emphasize the dependence of optimal posteriors from the marginal distribution.

Following Fosgerau et al. (2023) we show that problem (A.1) is concave. Without abuse of notation, we omit V in the argument and denote the objective function as U. The derivative equals to

$$U'(q) = -p^{H}(q) + \lambda c(p^{H}(q)) + (V - \lambda c(p^{L}(q))) + (1 - q)((p^{H}(q))' - \lambda c'(p^{H}(q))(p^{H}(q))') - (q)\lambda c'(p^{L}(q))(p^{L}(q))'$$

where $(p^H(q))'$ and $(p^L(q))'$ are the derivatives of the posteriors with respect to q. Using optimality condition $\lambda c'(p^H) - \lambda c'(p^L) = 1$ and differentiating the Bayesian consistency constraint to get $(1-q)(p^H(q))' + q(p^L(q))' = p^L(q) - p^H(q)$ we obtain

$$U'(q) = p^{H}(q) - \lambda c(p^{H}(q)) - (V - \lambda c(p^{L}(q))) + \lambda c'(p^{L}(q))(p^{H}(q) - p^{L}(q)).$$
(A.4)

Differentiating the expression with respect to q one more time and using optimality condition for posteriors result in

$$U''(q) = \lambda c''(p^{L}(q))(p^{H}(q) - p^{L}(q))(p^{L}(q))'.$$

We show that inequality $(p^L(q))' < 0$ holds. Combining optimality condition for posteriors and differentiable Bayesian consistency condition we obtain that equality $(p^L(q))'\left(q\frac{c''(p^L(q))}{c''(p^H(q))} + 1 - q\right) = p^L(q) - p^H(q)$ holds. Thus $(p^L(q))' < 0$ holds and U''(q) < 0 holds and problem (A.1) is concave. Therefore, problem (A.1) has a unique solution that is determined from the first-order condition or on the boundary. However, because the inequality $p^H > \mu > p^L$ holds, the optimal q is interior and uniquely determined from the first-order condition.

If $V \notin (V_L, V_H)$ then the DM chooses degenerate distribution, optimal q equals to 0 if $V \leq V_L$ and equals to 1 if $V \geq V_H$.

To get comparative statics of the optimal q with respect to V, we employ the standard supermodularity argument: optimal interior q is decreasing in V because the mixed derivative of $\frac{\partial^2 U(q,V)}{\partial q \partial V} = -1$ is negative.

B Proofs of Section **3**

B.1 Intermediate results

Lemma B2. $p_V^L, p_V^H \in (0, 1).$

Proof. By a standard concavification argument, the following inequality holds:

$$1 - \lambda c'(p_V^H) > -\lambda c'(p_V^L).$$

Now, assume the contrary to what we want to prove, i.e., for a given V, the employer chooses the optimal signal π with the pair of posterior beliefs p_V^L, p_V^H such that at least one of them belongs to the boundary $\{0, 1\}$.

If $p_V^L = 0$, by (2), we have

$$1 - \lambda c'(p_V^H) > -\lambda \lim_{p \to 0^+} c'(p) > 1 - \lambda c'(\mu),$$

and therefore, because c is convex, it is the case that $\mu > p_V^H$. Likewise, if $p_V^H = 1$, again by (2), we obtain

$$-\lambda c'(p_V^L) < 1 - \lambda \lim_{p \to 1^-} c'(p) < -\lambda c'(\mu),$$

and therefore by convexity of c, it is the case that $\mu < p_V^L$. In either case, $p_V^L < \mu < p_V^H$ is violated, and the proof is complete.

Lemma B3. Both p_V^L and p_V^H are differentiable with respect to V in (V_L, V_H) .

Proof. Under Lemma 1, the concave closure of ϕ as defined in the static problem (2) for $p \in [p_V^L, p_V^H]$ is a straight line that is tangent to ϕ at p_V^L and p_V^H . This tangent is characterized by the following equality for $p \in [p_V^L, p_V^H]$

$$V - \lambda c(p_V^L) - \lambda c'(p_V^L)(p - p_V^L) = p_V^H - \lambda c(p_V^H) - [\lambda c'(p_V^H) - 1](p - p_V^H), \quad (B.1)$$

such that

$$\lambda c'(p_V^L) = \lambda c'(p_V^H) - 1. \tag{B.2}$$

By the strict convexity of c and (B.2), we can implicitly define p_V^H as a continuously differentiable function of p_V^L . Using this in (B.1), we have that

$$V = p_V^H - \lambda c(p_V^H) - [\lambda c'(p_V^H) - 1](p - p_V^H) - [-\lambda c(p_V^L) - \lambda c'(p_V^L)(p - p_V^L)].$$
(B.3)

Using (B.2) yields:

$$0 = -V + p_V^H - \lambda c(p_V^H) + \lambda c(p_V^L) + \lambda c'(p_V^L)(p_V^H - p_V^L).$$
(B.4)

Next, note that the RHS in (B.4) is a continuously differentiable function of p_V^L . Moreover its derivative with respect to p_V^L is given by

$$(1 - \lambda c'(p_V^H)) \cdot (p_V^H)' + \lambda c'(p_V^L) + \lambda c'(p_V^L) \cdot (p_V^H)' - \lambda c'(p_V^L) + \lambda c''(p_V^L)(p_V^H - p_V^L) = \lambda c''(p_V^L)(p_V^H - p_V^L) > 0,$$

where the last equality comes from (B.2). The implicit function theorem implies that p_V^L is a continuously differentiable function of V and, consequently, so is p_V^H .

Lemma B4. Both p_V^L and p_V^H are increasing with respect to V in (V_L, V_H) .

Proof. Differentiating tangent optimality conditions (B.2) and (B.4) with respect to V gives the system

$$\frac{\partial p_V^L}{\partial V} = \frac{1}{\lambda c''(p_V^L)(p_V^H - p_V^L)},\tag{B.5}$$

$$\frac{\partial p_V^H}{\partial V} = \frac{1}{\lambda c''(p_V^H)(p_V^H - p_V^L)}.$$
(B.6)

By convexity of c, together with inequality $p_V^H > p_V^L$, both derivative are positive. \Box

B.2 Proof of Lemma 2

For every $V \in (V_L, V_H)$, the optimal signal π_V assigns to the two respective posteriors, p_V^L and p_V^H , probability

$$\pi_V(p_V^L) = \frac{p_V^H - \mu}{p_V^H - p_V^L} \text{ and } \pi_V(p_V^H) = \frac{\mu - p_V^L}{p_V^H - p_V^L},$$

and the employer's indirect expected utility in (V_L, V_H) is

$$g(V) = \pi_V(p_V^H) \left(p_V^H - \lambda c(p_V^H) \right) + \pi_V(p_V^L) \left(V - \lambda c(p_V^L) \right).$$

Since p_V^L and p_V^H are differentiable in V, so is g. By the Envelope Theorem, we have

$$g'(V) = \pi_V(p_V^L) > 0.$$

Thus, g is strictly increasing in (V_L, V_H) . Then, simple algebra yields

$$\frac{\partial \pi_V(p_V^L)}{\partial V} = \frac{\partial p_V^H}{\partial V}(\mu - p_V^L) + \frac{\partial p_V^L}{\partial V}(p_V^H - \mu),$$

which, by Lemma B4, is non-negative. Therefore, g is convex.

B.3 Proof of Lemma 3

Part 1 follows directly from (8) combined with Lemma 2.

By definition we have $V_T = 0$. Then, Part 2 follows directly from the fact that $g(V) \in (V_L, V_H)$ for all $V \in [0, V_H)$.

B.4 Proof of Theorem 1

The proof follows directly from Lemmas 1 and 3.

In particular, by $V_T = 0$, we get $V_T < V_L$. Hence, $\operatorname{supp}(\pi_{V_T}) = \{\mu\}$ and $\alpha_T(\mu) = 1$. Moreover, for every $i \in \{1, \ldots, T-1\}$, we have $V_L < V_i < V_H$, and therefore $\operatorname{supp}(\pi_{V_i}) = \{p_{V_i}^L, p_{V_i}^H\}$ with $\alpha_i(p_{V_i}^L) = 0$ and $\alpha_i(p_{V_i}^H) = 1$.

C Proofs of Section 4

C.1 Proof of Proposition 1

(i) \iff (ii): For each $s \in \{H, L\}$, we have:

$$p_i^s = \frac{\mu}{\mu + (1-\mu)\frac{\sigma_i(s|B)}{\sigma_i(s|G)}} > p_i^s = \frac{\mu}{\mu + (1-\mu)\frac{\sigma_j(s|B)}{\sigma_j(s|G)}} \iff \frac{\sigma_i(s\mid G)}{\sigma_i(s\mid B)} > \frac{\sigma_j(s\mid G)}{\sigma_j(s\mid B)}.$$

(i) \Rightarrow (iii): By the Bayes rule the passing probability under test k equals to $(\mu - p_k^L)/(p_k^H - p_k^L)$. The required follows from the inequalities

$$\frac{\mu - p_j^L}{p_j^H - p_j^L} > \frac{\mu - p_j^L}{p_i^H - p_j^L} < \frac{\mu - p_i^L}{p_i^H - p_i^L}.$$

(iii) \Rightarrow (i): Let the passing probability under test *i* be lower than under test *j*, but signal π_i be not more difficult than π_j . In this case, at least one inequality $p_j^H \ge p_i^H, p_j^L \ge p_i^L$ holds. In the first case for a candidate $\mu = p_i^H$ and in the second case for a candidate $\mu = p_j^L$, the passing probability of test *i* is higher. Thus, the signal π_i has to be more difficult than π_j .

(i) \Rightarrow (iv): Applying the Bayes rule and using (iii) implies:

$$\frac{\sigma_i(H|B)}{\sigma_j(H|B)} = \frac{\sigma_i(B|H)\sigma_i(H)}{\sigma_j(B|H)\sigma_j(H)} < 1.$$

Similarly,

$$\frac{\sigma_i(L|G)}{\sigma_j(L|G)} > 1.$$

(iv) \Rightarrow (i): Obviously, the inequality $\sigma_i(L|G) > \sigma_j(L|G)$ implies $\sigma_i(H|G) < \sigma_j(H|G)$. And, since $\sigma_i(H|B) < \sigma_j(H|B)$, (iv) is satisfied.

C.2 Proof of Theorem 2

For the first part, it is sufficient to show that both optimal posterior beliefs in the static problem (7) are increasing functions of the outside option. The latter follows from the proof of Lemma 2.

For the second part, note that by Proposition 1 and Lemma 1 continuation value

$$\lim_{T \to \infty} p_{V_1}^H = V_H.$$

In the solution to problem (7) with an outside option V_H , the solution to the firstorder conditions from Lemma 1 implies that the lower optimal posterior equals the prior. Therefore, by the continuity

$$\lim_{T \to \infty} p_{V_1}^L = \mu.$$

Finally, by Theorem 1, recall that the probability of the first candidate being hired is $(\mu - p_{V_1}^L)/(p_{v_1}^H - p_{V_1}^L)$, and therefore statements (a) and (b) follow trivially.

D Proofs of Section 5

D.1 Proof of Proposition 2

Proof. We rewrite the first-order optimality conditions for posterior beliefs p^L and p^H as identity for Bregman divergences. Adding and subtracting V in the right-hand side of the brackets of the equality

$$p^H - \lambda c(p^H) - (V - \lambda c(p^L)) = -\lambda c'(p^L)(p^H - p^L),$$

and using condition

$$\lambda c'(p^L) + 1 = \lambda c'(p^H),$$

we get that

$$c(V) - c(p^{H}) - c'(p^{H})(V - p^{H}) = c(V) - c(p^{L}) - c'(p^{L})(V - p^{L})$$
(D.1)

holds. Given that function c is strictly convex, we observe that identity (D.1) is the equality between two Bregman divergences

$$D(\mu, p^H) = D(\mu, p^L), \qquad (D.2)$$

where D(x, y) = c(x) - c(y) - c'(y)(x - y) is a Bregman divergence.⁷ Additionally, we substitute $\mu = V$. To analyze the equality (D.3) it is convenient to identify the test (p^L, p^H) as a pair (Δ^L, Δ^H) , where $\Delta^L = \mu - p^L$ and $\Delta^H = p^H - \mu$. Clearly, the first candidate is favored if and only if $\Delta^L > \Delta^H$ holds. With new notation, we can write

$$D(\mu, \mu + \Delta^H) = D(\mu, \mu - \Delta^L).$$
(D.3)

Taking arbitrary $\Delta \in \mathbb{R}$ and using the first fundamental theorem of calculus gives

$$D(\mu, \mu + \Delta) = \int_0^\Delta \frac{\partial D(\mu, \mu + \delta)}{\partial \delta} d\delta = \int_0^\Delta \delta c''(\mu + \delta) d\delta$$
(D.4)

and

$$D(\mu, \mu - \Delta) = \int_0^\Delta \frac{\partial D(\mu, \mu - \delta)}{\partial \delta} d\delta = \int_0^\Delta \delta c''(\mu - \delta) d\delta.$$

Conditions $c \in C_1$ or $c \in C_2$ pin down the inequality between $c''(\mu + \delta)$ and $c''(\mu - \delta)$ for each $\delta > 0$. In particular, let $c \in C_1$. If $\mu > z$ then $c''(\mu + \delta) > c''(\mu - \delta)$ for $\delta > 0$ and, therefore, for applying (D.4) and (D.4) to identity (D.3), it must be that $\Delta^L > \Delta^H$. If $\mu = z$ then $c''(\mu + \delta) = c''(\mu - \delta)$ for $\delta > 0$ and, it must be that $\Delta^L = \Delta^H$. If $\mu < z$ then $c''(\mu + \delta) < c''(\mu - \delta)$ for $\delta > 0$ and, it must be that $\Delta^L > \Delta^H$.

Therefore, if $c \in C_1$ then the first candidate is favored if $\mu > z$, the second candidate is favored if $\mu < z$, and no candidate is favored if $\mu = z$. The cases $c \in C_2$ and $c \in C_3$ are analogous.

⁷See for example, Hébert and Woodford (2023) for the application of Bregman divergences to the rational inattention problems.

D.2 Proof of Proposition 3

Simple algebra shows that the system has unique solution

$$p_V^L = V - \frac{1}{4\lambda}, p_V^H = V + \frac{1}{4\lambda}.$$
 (D.5)

Substituting the solution gives the value of the problem as

$$g(V) = V + \lambda(\mu - V + \frac{1}{4\lambda})^2$$

We show that for a given T > 2 inequality $q_{1T} > q_{2T}$ holds. Using the derived expressions above, we get that

$$q_{1T} = 2\lambda \Big(\mu - V_1 + \frac{1}{4\lambda}\Big); \qquad q_{2T} = 4\lambda^2 \Big(V_1 - \mu + \frac{1}{4\lambda}\Big) \Big(\mu - V_2 + \frac{1}{4\lambda}\Big),$$

moreover equality

$$V_1 = V_2 + \lambda \left(\mu - V_2 + \frac{1}{4\lambda}\right)^2$$

holds. We denote $t = \mu - V_2 + \frac{1}{4\lambda}$, thus,

$$q_{1T} = 2\lambda(t - \lambda t^2),$$

$$q_{2T} = 4\lambda^2 t \left(\frac{1}{2\lambda} - t + \lambda t^2\right).$$

Therefore, the inequality $q_{1T} > q_{2T}$ is equivalent to the inequality $\frac{1}{2\lambda} > t$. Because inequality $V_2 > \mu - \frac{1}{4\lambda}$ holds for all T > 2, inequality $q_{1T} > q_{2T}$ also holds.

By the Bayes rule the inequality $q_{tT} > q_{t+1T}$ is equivalent to the inequality $q_{1T-t+1} > q_{2T-t+1}$, therefore, the inequality holds.

Finally, if T = 2 then $q_{12} = 2\lambda(\mu - V_1 + \frac{1}{4\lambda}) = \frac{1}{2}$ because in this case $V_1 = \mu$. Therefore, $q_{T-1T} = q_{TT}$ for all T.

D.3 Proof of Theorem 3

Instead of analyzing the general dynamic problem (4) - (6), it is convenient to consider a collection of static problems (see Lemma A1) that are introduced in Appendix A. Additionally, it is convenient to use slightly different notations for the proof. We denote s_T and R_T to be the optimal passing probability and the optimal attained value in the problem with T available candidates correspondingly⁸

$$s_T = \underset{s \in [0,1]}{\arg \max} U(s, R_{T-1}),$$

 $R_T = U(s_T, R_{T-1}).$

The proof consists of four parts. First, we show that for large enough T approximation $\frac{1}{s_T} - \frac{1}{s_{T-1}} \approx \frac{1}{2}$ holds. Second, we reformulate the statement of the theorem

⁸To emphasize the difference, in Appendix A in the DM's problem r_i and V_i are passing probability and continuation value for different *i* and fixed *T*. In this section, we vary *T* instead.

into the convergent property of the sequence. Third, using step 1, we simplify the sequence from step 2. Fourth, using step 1 again, we show the desired property of the sequence and conclude the proof.

Part 1.

We consider the limit of expression $\frac{1}{s_T} - \frac{1}{s_{T-1}}$ when $T \to \infty$.

From Lemma A1 function U(s, R) is strictly concave in s, therefore, there exists a function $v : \mathbb{R} \to \mathbb{R}$ such that $v(R_{T-1}) = s_T$ holds. Moreover, from the Implicit function theorem function v(.) is differentiable and from Lemma A1 equality $v'(R_{T-1}) = \frac{1}{U_{ss}(s_T, R_{T-1})}$ holds. Denoting $\varphi(x) = \frac{1}{v(x)}$ by the mean value theorem we get that

$$\frac{1}{s_T} - \frac{1}{s_{T-1}} = \varphi(g(R_{T-1})) - \varphi(R_{T-1}) = (g(R_{T-1}) - R_{T-1})\varphi'(\xi), \quad (D.6)$$

where $\xi \in (R_{T-1}, g(R_{T-1}))$. By the chain rule $\varphi'(x) = -\left(\frac{1}{v(x)}\right)^2 v'(x)$ holds. Using the expression for $v'(R_{T-1})$ we can write

$$\frac{1}{s_T} - \frac{1}{s_{T-1}} = -(g(R_{T-1}) - R_{T-1}) \left(\frac{1}{v(\xi)}\right)^2 \frac{1}{U_{ss}(v(\xi),\xi)}$$

We approximate expression $(g(R_{T-1}) - R_{T-1})$ using the Taylor expansion. Clearly, identities $U(s_T, R_{T-1}) = g(R_{T-1})$ and $U(0, R_{T-1}) = R_{T-1}$ hold. We use Taylor expansion for value $U(s_T, R_{T-1})$:

$$U(s_T, R_{T-1}) = U(0, R_{T-1}) + s_T \times U_s(0, R_{T-1}) + \frac{1}{2}s_T^2 \times U_{ss}(\chi, R_{T-1}),$$

where $\chi \in (0, s_T)$. Therefore, equality

$$g(R_{T-1}) - R_{T-1} = s_T \times U_s(0, R_{T-1}) + \frac{1}{2}s_T^2 \times U_{ss}(\chi, R_{T-1})$$
(D.7)

holds. Using mean value theorem for function $U_s(s, R_{T-1})$ at point $s = s_T$ we get

$$U_s(s_T, R_{T-1}) = U_s(0, R_{T-1}) + s_T \times U_{ss}(\psi, R_{T-1}),$$

where $\psi \in (0, s_{T-1})$. From the first-order condition for s_T we get equality $U_s(s_T, R_{T-1}) =$ 0, therefore, equality

$$U_s(0, R_{T-1}) = -s_T \times U_{ss}(\psi, R_{T-1})$$
(D.8)

holds. Substituting expression (D.8) into (D.7) results in

$$g(R_{T-1}) - R_{T-1} = -s_T^2 \times U_{ss}(\varphi, R_{T-1}) + \frac{1}{2}s_T^2 \times U_{ss}(\chi, R_{T-1}).$$
(D.9)

Therefore, we can write

$$\frac{1}{s_T} - \frac{1}{s_{T-1}} = \frac{s_T^2 \times U_{ss}(\varphi, R_{T-1}) - \frac{1}{2}s_T^2 \times U_{ss}(\chi, R_{T-1})}{v(\xi)^2 U_{ss}(v(\xi), \xi)}.$$
 (D.10)

From the mean value theorem $\xi \in (R_{T-1}, R_T)$ holds, therefore, we observe that $\lim_{T\to\infty} \frac{s_T}{v(\xi)} = 1$ because sequence $(s_T)_T$ converges. Similarly, $\lim_{T\to\infty} \frac{U_{ss}(\varphi, R_T)}{U_{ss}(v(\xi),\xi)} = 1$ and $\lim_{T\to\infty} \frac{U_{ss}(\chi, R_T)}{U_{ss}(v(\xi),\xi)} = 1$ holds. The sequence $(R_T)_T$ converges to V_H therefore the values of the second derivative converge to a value $U_{ss}(0, V_H)$, which equals to $U_{ss}(0, V_H) = -\lambda c''(\mu)(V_H - \mu)^2$ from the proof of Lemma A1. By combining obtained limiting behaviors, we get that

$$\lim_{T \to \infty} \frac{s_T^2 \times U_{ss}(\varphi, R_T) - \frac{1}{2}s_T^2 \times U_{ss}(\chi, R_T)}{h(\xi)^2 U_{ss}(h(\xi), \xi)} = \frac{1}{2}$$
(D.11)

and, therefore,

$$\lim_{T \to \infty} \left(\frac{1}{s_T} - \frac{1}{s_{T-1}} \right) = \frac{1}{2}.$$

Part 2.

First, we show that exists a number \overline{T} such that inequality $q_{1T} > q_{iT}$ holds for all $i \in [2, T - \overline{T}' + 1]$ for all $T > \overline{T}$. The statement implies that if there are at least $\overline{T} + 1$ candidates, then the first candidate has the largest probability of being hired amount first $T - \overline{T}' + 1$ candidates.

To prove the statement, it is sufficient to show that for $T > \overline{T}$ inequality $q_{1T} > q_{2T}$ holds. If $q_{1T} > q_{2T}$ holds for some T then $q_{2,T+1} > q_{3,T+1}$ simply by the Bayes rule. Therefore, inequalities $q_{1,T+1} > q_{2,T+1} > q_{3,T+1}$ hold. Continuing the argument we get that if $q_{1T} > q_{2T}$ holds for all $T > \overline{T}$ then inequalities

$$q_{1T} > q_{2T} > \ldots > q_{T-\bar{T}'+1,T}$$

hold, and the first candidate has the largest probability of being hired amount first $T - \overline{T}' + 1$ candidates.

We show inequality $q_{1T} > q_{2T}$ using the results from Part 1. We rewrite $q_{1T} > q_{2T}$ using passing probabilities:

$$q_{1T} > q_{2T} > \iff s_T > (1 - s_T)s_{T-1} \iff \frac{1}{s_T} - \frac{1}{s_{T-1}} < 1.$$

From Part 1, the last inequality holds for sufficiently large T.

Second, consider the case with \overline{T} available candidates. Let i_{\max} be the candidate with the largest probability of being hired and q_{\max} be the corresponding probability when there are \overline{T} available candidates. If $i_{\max} = 1$, then trivially, the statement from the beginning of this Part completes the proof of the theorem. If $i_{\max} \neq 1$, then we increase the number of available candidates and compare the probabilities of hiring a first candidate and candidate i_{\max} with an increased number of candidates. That is, let the number of candidates be $\widetilde{T} > \overline{T}$. The probability to hire the first candidate equals to $s_{\widetilde{T}}$ and the probability to hire candidate i_{\max} equals to $(1 - s_{\widetilde{T}}) \times (1 - s_{\widetilde{T}-1}) \times \ldots \times (1 - s_{\overline{T}+1}) \times q_{\max}$.

Therefore, to complete the proof of the theorem, it is sufficient to show that

$$\lim_{\tilde{T}\to\infty} \frac{s_{\tilde{T}}}{(1-s_{\tilde{T}})\times(1-s_{\tilde{T}-1})\times\ldots\times(1-s_{\tilde{T}+1})} = \infty.$$
 (D.12)

If the limiting behavior (D.12) holds, then exists T' such that if there are T', available candidates, then the probability of hiring the first candidate is larger than the candidate i_{max} .

Part 3.

We rewrite the expression under the limit (D.12) as

$$\frac{s_{\tilde{T}}}{(1-s_{\tilde{T}})\times(1-s_{\tilde{T}-1})\times\ldots\times(1-s_{\bar{T}+1})} = \frac{s_{\tilde{T}}}{s_{\tilde{T}-1}(1-s_{\tilde{T}})} \times \frac{s_{\tilde{T}-1}}{s_{\tilde{T}-2}(1-s_{\tilde{T}-1})} \times \ldots \times \frac{s_{\bar{T}+2}}{s_{\bar{T}+1}(1-s_{\bar{T}+2})} \times s_{\bar{T}+1} \times \frac{1}{1-s_{\bar{T}+1}}.$$

We denote $y_k = \frac{s_{\bar{T}+k+1}}{s_{\bar{T}+k}(1-s_{\bar{T}+k+1})}$ for $k \ge 1$. Therefore, to show limiting behavior (D.12) is equivalent to show

$$\lim_{K \to \infty} \prod_{k=1}^{K} y_k = \infty.$$
 (D.13)

Using the results from Part 1, we express the value y_k . We use the approximation $\frac{1}{s_{\tilde{T}+k+1}} - \frac{1}{s_{\tilde{T}+k}} \approx \frac{1}{2}$.⁹ Routine algebra shows that

$$\frac{s_{\tilde{T}+k+1}}{s_{\tilde{T}+k}(1-s_{\tilde{T}+k+1})} \approx 1 + \frac{1}{2} \frac{s_{\tilde{T}+k+1}}{1-s_{\tilde{T}+k+1}}.$$

Using Taylor approximation of function $f(x) = \frac{1}{1-x} \approx 1 + x$ around x = 0 gives the first order approximation of y_k as

$$y_k \approx 1 + \frac{1}{2} s_{\tilde{T}+k+1}.$$
 (D.14)

To show (D.13) is equivalent to show

$$\lim_{K \to \infty} \log \prod_{k=1}^{K} y_k = \infty$$

We simplify the expression under the limit and use the derived approximation (D.14)

$$\log \prod_{k=1}^{K} y_k = \sum_{k=1}^{K} \log y_k \approx \sum_{k=1}^{K} \log(1 + \frac{1}{2}s_{\tilde{T}+k+1}) \approx \frac{1}{2} \sum_{k=1}^{K} s_{\tilde{T}+k+1}$$

Therefore, to show the limiting behavior (D.12) we need to show that

$$\lim_{K \to \infty} \sum_{k=1}^{K} s_{\tilde{T}+k+1} = \infty.$$
 (D.15)

holds.

 $^{^{9}}$ To show that the sequence in (D.12) diverges is sufficient to consider zero-order approximation. If the sequence diverges, the higher-order terms of the infinitesimal sequence do not change the limiting behavior.

Part 4.

For convenience, we denote $l = \tilde{T} + k + 1$. Therefore, the limit (D.12) is equivalent to

$$\lim_{L \to \infty} \sum_{l=1}^{L} s_l = \infty.$$
 (D.16)

Using again approximation $\frac{1}{s_{l+1}} - \frac{1}{s_l} \approx \frac{1}{2}$ we observe that the sequence $(\frac{1}{s_l})_l$ grows linearly in l. Therefore, the sequence $(s_l)_l$ grows as $\frac{1}{l}$ in l. Thus the sum $\sum_{l=1}^{L} s_l$ grows as the sum of harmonic sequence $(\frac{1}{l})_l$. The sum of the harmonic sequence divergences; therefore, the limiting behavior (D.16) holds. Therefore, by Part 2, T'exists, such that if the number of candidates is larger than T', the probability of hiring the first candidate is the largest.

E Proofs of Section 6

E.1 Proof of Theorem 4

QUADRATIC COST: Let us first take any two $\mu_1, \mu_2 \in M$, which are sufficiently close to each other, i.e.,

$$\mu_1 \in \left(p_{\mu_2}^L, p_{\mu_2}^H\right) \text{ and } \mu_2 \in \left(p_{\mu_1}^L, p_{\mu_1}^H\right).$$

Then, using the conventional notation $\kappa := 1/4\lambda$, by (D.5) we obtain

$$U_{1} := \frac{\mu_{1} - p_{\mu_{2}}^{L}}{p_{\mu_{2}}^{H} - p_{\mu_{2}}^{L}} p_{\mu_{2}}^{H} + \frac{p_{\mu_{2}}^{H} - \mu_{1}}{p_{\mu_{2}}^{H} - p_{\mu_{2}}^{L}} \mu_{2}$$
$$= \frac{\mu_{1} - \mu_{2} + \kappa}{2\kappa} (\mu_{2} + \kappa) + \frac{\mu_{2} - \mu_{1} + \kappa}{2\kappa} \mu_{2}$$
$$= \frac{\mu_{1} + \mu_{2} + \kappa}{2}.$$

Then, by symmetry, we obviously get $U_1 = U_2$.

Then, suppose that $\mu_1 < \mu_2$ are not sufficiently close to each other, i.e., let $\mu_2 > p_{\mu_1}^H$. Now, let us introduce the convenient notation:

$$\begin{aligned} h_L(V) &:= p_V^L, \\ h_H(V) &:= p_V^H, \end{aligned}$$

and then take the useful composite function

$$h := h_L \circ h_H. \tag{E.1}$$

Then, notice that for every $V \in (0, 1)$, we trivially obtain

$$h(V) = V$$

Hence, it will necessarily be the case that $\mu_1 < p_{\mu_2}^L$, meaning that regardless who arrives first, candidate 2 will be hired without any of the two being interviewed. Hence, we will have $U_1 = U_2$.

ENTROPIC COST: We will now use the notational convention $\kappa := 1/\lambda$, and we will use again the function h as defined in (E.1). Then, using the formulas for the posterior beliefs,

$$p_{\mu}^{L} = \frac{e^{\mu/\lambda} - 1}{e^{1/\lambda} - 1} \qquad \qquad p_{\mu}^{H} = \frac{e^{1/\lambda} - e^{(1-\mu)/\lambda}}{e^{1/\lambda} - 1}, \qquad (E.2)$$

we obtain

$$h'(V) = \frac{\kappa e^{\kappa h_H(V)}}{e^{\kappa} - 1} \cdot \frac{\kappa e^{\kappa(1-V)}}{e^{\kappa} - 1},$$

which is obviously positive. Furthermore, we have

$$h''(V) = \frac{\kappa e^{\kappa(1-V)}}{e^{\kappa} - 1} - 1.$$

which has a unique root

$$\mu_0 := 1 - \frac{1}{\kappa} \log \frac{e^{\kappa} - 1}{\kappa} \in (0, 1).$$

In addition we have h''(V) > 0 if and only if $V < \mu_0$. This means that h is strictly convex below μ_0 and strictly concave above μ_0 . Therefore, it will be the case that

h(V) < V for all $V < \mu_0$, and h(V) > V for all $V > \mu_0$.

Part 1: Take an arbitrary $\mu_1 < \mu_0$. By h_L being continuously increasing, for each $\mu_2 > p_{\mu_1}^H$ there is some $\varepsilon > 0$ such that

$$p_{\mu_2}^L = h_L(\mu_2) = h_L(p_{\mu_1}^H) + \varepsilon = h(\mu_1) + \varepsilon.$$

By taking μ_2 sufficiently close to $p_{\mu_1}^H$, it will be the case that the corresponding ε will satisfy

$$\varepsilon < \mu_1 - h(\mu_1).$$

By strict convexity of h below μ_0 , the righthand side is strictly positive. As a result, we obtain

$$p_{\mu_2}^L < \mu_1.$$

Thus, by $\mu_2 > p_{\mu_1}^H$, if candidate 2 goes first, he will be directly hired without any interview, and therefore we will have $U_2 = \mu_2$. On the other hand, by $p_{\mu_2}^L < \mu_1$, if candidate 1 goes first, there will be an informative interview, meaning that it is not optimal to hire candidate 2 without an interview, and therefore $U_1 > \mu_2$. As a result, we conclude that $U_1 > U_2$.

Part 2: We follow same steps as in the previous part. Take an arbitrary $\mu_1 > \mu_0$. Then, there is some $\mu_2 < p_{\mu_1}^H$ and some $\varepsilon < h(\mu_1) - \mu_1$ such that $p_{\mu_2}^L = h(\mu_1) - \varepsilon$. Hence, it will be the case that

$$p_{\mu_2}^L > \mu_1.$$

Then, by the same argument as above, we have $U_1 = \mu_2$ and $U_2 > \mu_2$, which implies $U_1 < U_2$, thus completing the proof.

F Proofs of Section 7

F.1 Proof of Proposition 5

We assume that the solution to the restricted problem is interior and later show that it is indeed optimal. Optimal interior q^* solves

$$\frac{\partial U}{\partial r}(q^*, V_1(q^*, \mu)) + \frac{\partial U}{\partial V}(q^*, V_1(q^*, \mu)) \times \frac{dU}{dr}(q^*, V_2(q^*, \mu)) = 0.$$

Using the chain rule, we can express this first-order condition as

$$\frac{\partial U}{\partial r}(q^*, V_1(q^*, \mu)) + \sum_{j=2}^{T-1} \left(\frac{\partial U}{\partial r}(q^*, V_j(q^*, \mu)) \prod_{k=1}^{j-1} \frac{\partial U}{\partial V}(q^*, V_k(q^*, \mu)) \right) = 0.$$
(F.1)

We will show that the first term in the first-order condition is negative, that inequality $\frac{\partial U}{\partial r}(q^*, V_1(q^*, \mu)) < 0$ holds. We show this fact by the contradiction, analyzing the second term in the first-order condition.

First, we observe that for general values r, V, the following partial derivatives are equal to

$$\frac{\partial U}{\partial V}(r,V) = 1 - r, \quad \frac{\partial^2 U}{\partial V \partial r}(r,V) = -1,$$

therefore, partial derivative $\frac{\partial U}{\partial V}(r, V)$ is always positive and function $\frac{\partial U}{\partial r}(r, V)$ is decreasing in V.

Second, let inequality $\frac{\partial U}{\partial r}(q^*, V_1(q^*, \mu)) \geq 0$ hold. The partial derivative of function in r is decreasing in the second argument, therefore, inequality

 $\frac{\partial U}{\partial r}(q^*, V_{i+1}(q^*, \mu)) > 0 \text{ hold for all } i > 1. \text{ Thus, the second term in the first-order condition (F.1) is positive. Therefore, the left-hand side of the equation (F.1) is positive. We reach a contradiction, thus, inequality <math>\frac{\partial U}{\partial r}(q^*, V_1(q^*, \mu)) < 0$ holds. To compare q^* with q^{**} we observe that inequality $\frac{\partial U}{\partial r}(q^*, \mu) > 0$ holds. Indeed, if

To compare q^* with q^{**} we observe that inequality $\frac{\partial U}{\partial r}(q^*,\mu) > 0$ holds. Indeed, if this inequality does not hold, then the argument from the above paragraph suggests that the left-hand side of the first-order condition (F.1) is negative. Optimal q^{**} in the T = 2 case solves $\frac{\partial U}{\partial r}(q^{**},\mu) = 0$. Function $\frac{\partial U}{\partial r}(r,V)$ is decreasing in r, therefore, inequality $q^{**} > q^*$ holds.

To compare q^* with q^{***} we consider the auxiliary static problem with an outside option $V_1(q^*, \mu)$:

$$\max\{U(r, V_1(q^*, \mu))\}$$

In this problem, the value of the outside option equals the obtained outside option value in the restricted problem. Let q^{****} be a solution to this problem, therefore q^{****} solves $\frac{\partial U}{\partial r}(q^{****}, V_1(q^*, \mu)) = 0$. Function $\frac{\partial U}{\partial r}(r, V)$ is decreasing in r, therefore the inequality $q^* > q^{****}$ holds.

The unrestricted problem is a static problem with an outside option V_1 . Clearly, the inequality $V_1 > V_1(q^*, \mu)$ holds. Optimal r in the static problem decreases in the outside option value V, therefore, inequality $q^{****} > q^{***}$ holds and by transitivity inequality $q^* > q^{***}$ also holds.

We show that the optimal q^* is interior by taking derivatives of the function $U(r, V_1(r, \mu))$ on the boundary. We consider the case r = 0, and the analysis of case

r = 1 is identical. We observe that $V_i(0, \mu) = \mu$ for all i < T. Additionally, equality $\frac{\partial U}{\partial V}(0, V_i(0, \mu)) = 1$ holds. Therefore, using the left-hand side of the expression (F.1) we obtain that

$$\left. \frac{\partial U}{\partial r}(r, V_1(r, \mu)) \right|_{r=0+0} = (T-1) \frac{\partial U}{\partial r}(0, \mu) > 0.$$

The last inequality holds because the static problem with an outside option μ has an interior solution and by Lemma A1 function $U(r, \mu)$ is concave in r.

References

- Anderson, S. and Renault, R. (2018). Firm pricing with consumer search. Handbook of Game Theory and Industrial Organization, II.
- Arieli, I., Babichenko, Y., Smorodinsky, R., and Yamashita, T. (2023). Optimal persuasion via bi-pooling. *Theoretical Economics*, 18(1):15–36.
- Armstrong, M. (2017). Ordered consumer search. Journal of the European Economic Association, 15(5):989–1024.
- Bartoš, V., Bauer, M., Chytilová, J., and Matějka, F. (2016). Attention discrimination: Theory and field experiments with monitoring information acquisition. *American Economic Review*, 106(6):1437–1475.
- Bertheau, A., Larsen, B., and Zhao, Z. (2023). What makes hiring difficult? Evidence from linked survey-administrative data.
- Bertrand, M. and Duflo, E. (2017). Field experiments on discrimination. *Handbook* of Economic Field Experiments, I.
- Beyhaghi, H. and Cai, L. (2023). Recent developments in pandora box problem: Variants and applications.
- Caplin, A. and Dean, M. (2013). Behavioral implications of rational inattention with Shannon entropy.
- Caplin, A., Dean, M., and Leahy, J. (2019). Rational inattention, optimal consideration sets, and stochastic choice. *The Review of Economic Studies*, 86(3):1061– 1094.
- Caplin, A., Dean, M., and Leahy, J. (2022). Rationally inattentive behavior: Characterizing and generalizing Shannon entropy. *Journal of Political Economy*, 130(6):1676–1715.
- Chade, H., Eeckhout, J., and Smith, L. (2017). Sorting through search and matching models in economics. *Journal of Economic Literature*, 55(2):493–544.
- Charness, G., Oprea, R., and Yuksel, S. (2021). How do people choose between biased information sources? evidence from a laboratory experiment. *Journal of* the European Economic Association, 19(3).
- Correa, J., Cristi, A., Feuilloley, L., and Oosterwijk, Timand Tsigonias-Dimitriadis, A. (2024). The secretary problem with independent sampling. *Management Sci*ence.
- Cotterell, T. (2019). A little biased? interview order does matter. Forbes. Available at: https://www.forbes.com/councils/forbeshumanresourcescouncil/ 2019/03/22/a-little-biased-interview-order-does-matter/ (Accessed: 2024-12-09).

- Dean, M. and Neligh, N. (2023). Experimental tests of rational inattention. Journal of Political Economy, 131(12):3415–3461.
- Denti, T. (2022). Posterior separable cost of information. American Economic Review, 112(10):3215–3259.
- Doval, L. (2018). Whether or not to open pandora's box. *Journal of Economic Theory*, 175:127–158.
- Fosgerau, M., Sethi, R., and Weibull, J. (2023). Equilibrium screening and categorical inequality. *American Economic Journal: Microeconomics*, 15(3):201–242.
- Gentzkow, M., Shapiro, J., and Stone, D. (2014). Media bias in the marketplace: Theory.
- Georgiadis-Harris, A. (2024). Preparing to act. Technical report, Working Paper.
- Gill, D. and Sgroi, D. (2012). The optimal choice of pre-launch reviewer. *Journal* of *Economic Theory*, 147:1247–1260.
- Hébert, B. and Woodford, M. (2021). Neighborhood-based information costs. American Economic Review, 111(10):3225–3255.
- Hébert, B. and Woodford, M. (2023). Rational inattention when decisions take time. Journal of Economic Theory, 208:105612.
- Jain, V. and Whitmeyer, M. (2021). Search and competition with endogenous investigations. Technical report, Working Paper.
- Kamenica, E. and Gentzkow, M. (2011). Bayesian persuasion. American Economic Review, 101(6):2590–2615.
- Kim, J. H. J., Kim, K., and Pease, M. (2022). Unemployment duration under flexible information acquisition.
- Laker. В. (2023).Understanding the impact of inter-Available scheduling Forbes. view on hiring decisions. at: https://www.forbes.com/sites/benjaminlaker/2023/12/14/ understanding-the-impact-of-interview-scheduling-on-hiring-decisions/ (Accessed: 2024-12-09).
- Matějka, F. and McKay, A. (2015). Rational inattention to discrete choices: A new foundation for the multinomial logit model. *American Economic Review*, 105(1):272–298.
- McCall, J. J. (1970). Economics of information and job search. *The Quarterly Journal of Economics*, 84(1):113–126.
- Mensch, J. and Malik, K. (2023). Posterior-mean separable costs of information acquisition. arXiv preprint arXiv:2311.09496.

- Miao, J. and Xing, H. (2024). Dynamic discrete choice under rational inattention. *Economic Theory*, 77(3):597–652.
- Morris, S. and Strack, P. (2019). The Wald problem and the relation of sequential sampling and ex ante information costs.
- Mortensen, D. and Pissarides, C. (1999). New developments on models of search in the labor market. *Handbook of Labor Economics*, III.
- Neumark, D. (2018). Experimental research on labor market discrimination. *Journal* of *Economic Literature*, 56:799–866.
- Salop, S. (1973). Systematic job search and unemployment. Review of Economic Studies, 40:191–201.
- Schlag, K. and Zapechelnyuk, A. (2024). Fair hiring procedures. Technical report, Working Paper.
- Selby Jennings (2023). Do employers interview the best candidate first? Available at: https://www.selbyjennings.com/blog/2023/04/ do-employers-interview-best-candidate-first (Accessed: 2024-12-09).
- Shaked, M. and Shanthikumar, G. (2007). Stochastic orders. Springer.
- Sims, C. (2003). Implications of rational inattention. Journal of Monetary Economics, 50(3):665–690.
- Steiner, J., Stewart, C., and Matejka, F. (2017). Rational inattention dynamics: Inertia and delay in decision-making. *Econometrica*, 85(2):521–553.
- Stigler, G. (1961). The economics of information. *Journal of Political Economy*, 69(3):213–225.
- Ursu, R. M., Wang, Q., and Chintagunta, P. K. (2020). Search duration. Marketing Science, 39(5):849–871.
- Weitzman, M. L. (1979). Optimal search for the best alternative. *Econometrica*, 47(3):641–654.
- Zhong, W. (2022). Optimal dynamic information acquisition. *Econometrica*, 90(4):1537–1582.