Secret Search*

June 24, 2019

Andrew John

Melbourne Business School, University of Melbourne, Australia

Ian King[1]

School of Economics, University of Queensland, Australia

Abstract

For high-profile positions, should applicant identities be made public within the organisation ("open search") or kept confidential ("secret search")? We construct a model where an organisation seeks to hire, but where candidates' abilities are private information unless it uses open search. Rejected applicants, under open search, suffer disutility. We find: salaries are lower under secret search, the expected ability of applicants decreases as the posted (open search) salary increases, secret search is preferred by organisations where quality of candidate is relatively unimportant, and organisations will, for some parameter values, choose secret search even when open search is more efficient.

JEL Codes: M51, M52

Key words: Executive recruitment, search, asymmetric information

Running Head: Secret Search

*Manuscript originally submitted 19 June 2017. Revision submitted 30 January 2019

[1]We thank three referees and the editor for their detailed and helpful comments. We would also like to thank Catherine de Fontenay, Begona Dominguez, Allen Head, Dmitri Kvasov, Benoit Julien, John Stachursky, Frank Staehler, Graham Voss, Jacob Wong, and seminar participants at the Australasian Economic Theory Workshop, the Economics Departments at Queens University, Tubingen University, the University of Adelaide, the University of Queensland, the University of Victoria, and at brown bag seminars in Economics Department and Melbourne Business School at the University of Melbourne. We are grateful to Timur Behlul for outstanding research assistance. The usual disclaimer applies. This project was partially funded by the Australian Research Council Discovery Project DP150103987.

> Selecting a CEO is an information-intensive decision. The information needed for selecting a CEO is usually highly particular and fine-grained. That is, it is usually information that can be derived only through direct experience in working with and observing the candidate. Some examples of particular factors about which information is needed include candidates' dispositions, working styles, and mannerisms; how they work with others; and, most importantly, whether they can really do what they say they are capable of doing. ... For outside candidates, this information is not so easily gathered. In particular, because of the confidentiality inherent in CEO searches, it is not possible to interview peers or subordinates for particular information about the candidate. [Khurana, 2005]

1 Introduction

When an organisation or firm tries to recruit someone for a high-profile position it must, as part of the process, make a decision about whether or not to make the identity of applicants public knowledge. That is, the organisation must choose whether to make the search "open" (known broadly across the organisation and, hence, typically outside of the organisation as well) or "secret" (known only to a small committee which is bound by confidentiality). Secret search is quite common, not only in the private sector but also in quasi-public institutions such as universities.[2] In this paper we analyse these search processes from the point of view of a profit-maximising firm that is aware of the implications of its choices on the expected pool of applicants and expected profits.

A key declared benefit of a secret search process is that more candidates are likely to apply if they do not face the potential embarrassment of being seen to be rejected from the position. That is, the chance of being observed as an unsuccessful applicant can deter applications, even of high-quality candidates. However, secrecy also has its costs. Most notably, it is more difficult to gather reliable information about candidates when the search process imposes confidentiality. This is the essence of the trade-off that we analyse: secret search increases the expected number of applications but decreases the amount of information available about each applicant.[3]

We model this tradeoff in an environment where a firm faces potential candidates whose ability levels are, initially, private information. The firm can learn the specific ability of an applicant by adopting an open search process. However, if the firm chooses such a process, applicants run the

[2] In academia, confidentiality is usually imposed in the early stages of the search for high-level administrators, but is often relaxed in the final stages; see Thuene and Tiede (2014) for some evidence on this.

[3] Secrecy, arguably, also has the cost of creating a higher probability of cronyism; we put that issue aside here.

risk of suffering a disutility from their public rejection if they are unsuccessful. When search is secret, candidates do not face this risk. The game has three stages. First, the firm announces its search strategy (open or secret) and posts a salary. Second, candidates choose whether or not to apply. Third, given the pool of available applicants, and the information available to it, the firm chooses an applicant to maximise equilibrium expected profits. We consider two variants. In the first, candidates abilities are drawn from a continuous uniform distribution. In the second variant of our model, candidates' abilities can be either high or low, with corresponding values of output for the firm.

Our models are thus very simple: there are only two parameters in the first variant (the maximum output premium of a high-quality candidate and the utility cost for an unsuccessful applicant) and only three in the second (the output premium of a high-quality candidate, the fraction of candidates that are high quality, and the utility cost). But the solutions are non-trivial. In both cases, the expected profit of a firm under open search depends in a complicated way upon the salary. In the second variant, the expected profit function is in fact non-monotonic: open search comes in two qualitatively different flavours, corresponding to different local maxima of the expected profit function. In one case the firm adopts a separating strategy in which only high-ability candidates apply. In the other, a pooling strategy is used, where both types of agents apply. We call these cases "open separating" and "open pooling", respectively. The implications for salaries are very different in the different cases.

Our first model provides some basic results on secret and open search. We find that, as intuition would suggest, a higher support of candidate abilities and a lower utility cost from rejection both favour secret search. We also find that the salary chosen under open search is inefficient. The second model allows us to understand how the *distribution* of agents affects the outcome of the model. As well as confirming (in a slightly different setting) key results from the first model, it reveals that the firm's choices about its hiring regime depend in interesting ways on the mix of agents in the candidate pool.

We use these models to consider the following questions. What features of the candidate pool cause a firm to choose secret, open separating, or open pooling search? What type of candidate is most likely to apply in these different cases? What are the implications for salary levels, and for the relationship between salary and productivity levels? What allocations would a (constrained) planner choose? When would these allocations differ from the equilibrium ones? Under a slightly different interpretation of the model, we can also ask: what type of firm would choose secret search?

Related Literature

The existing literature on our specific topic is very thin. We could not find a single article in the academic economics literature that directly addresses open versus secret search. Of course, the literature on search-theoretic models of the labour market is very well developed (see Rogerson et al. 2005 for a general survey and Wright et al. 2017 for a recent survey on directed search). Historically, this literature has focused primarily on candidates looking for jobs, rather than firms searching for workers. Some studies in the directed search literature have analysed recruitment (Julien et al. 2000, Albrecht et al. 2006, Galenianos and Kircher 2009, Wolthoff 2012, and Basov et al. 2014), but in these studies firms do not decide which search process to use. The mechanism design literature (in particular, McAfee 1993, Peters and Severinov 1997, Eeckhout and Kircher 2010, Guerrieri et al. 2010, Auster and Gottardi 2017, and Cai et al. 2017) considers choices of pricing processes, but does not allow firms to choose the processes considered in this paper.

In their survey of personnel economics, Oyer and Schaefer (2010) point out that there is significant heterogeneity in firm recruiting practices, and that this heterogeneity has been considerably under-researched. Going back to Rees (1966), there has long been a focus on the distinction between practices that affect the extensive margin (the number of applicants) and the intensive margin (the quantity of information per candidate); this distinction is a central theme of this paper. Another theme that we follow—namely how compensation can be used by firms to induce self-selection—was first examined in Salop and Salop (1976), although through an entirely different mechanism. Van Ours and Ridder (1992) determine that, unlike employee search, employer search tends to be simultaneous (as modelled here) rather than sequential.

Alonso (2014) emphasises that the hiring process, together with compensation, is an important determinant of a candidate's willingness to apply to a firm. He presents a model that is similar in spirit to ours, but with a very different focus. In both models applicants must receive a premium to offset application costs. In our framework, however, we focus on a particular cost that is incurred only by unsuccessful applicants in an open search process. Alonso, by contrast, examines the screening process that firms use. Villena-Rodin (2012) and Lester and Wolthoff (2016) consider models where firms pay a screening (or interview) cost per applicant (which, in principle with transferable utility, could be borne by either party in equilibrium). Lester and Wolthoff do consider the optimal stopping problem for the number of applicants to screen, but they explicitly rule out ("to keep the analysis as simple as possible") the type of comparison that is closest to the focus of the current study: random (costless) selection versus costly screening, and the consequences for the number of applications.

Our paper is specifically concerned with recruitment to senior positions. The management literature provides some research on the connections between CEO compensation and firm performance.

Some of this work asks which characteristics of CEOs are predictive of corporate performance (for example, Kaplan et al. 2012); some asks whether high compensation reflects reward for high performance, or simply rent extraction (for example, Bebchuk and Fried, 2003, and see in particular the survey by Edmans and Gabaix, 2016); some considers the role of intermediaries such as headhunting agencies (for example, Rajgopal et al. 2012). However, there again appears to be comparatively little research on the CEO *selection* process. Zajac (1990) is an exception, and was probably the first in this literature to argue that adverse selection matters for the choice of CEOs. Noting that the firm is likely to be better informed about internal candidates than external candidates, he hypothesised that internal candidates should be associated with better firm performance; consistent with this belief, he found that the return on assets was higher for firms that promoted from within. While we do not make a distinction between hiring from outside versus hiring from within the firm, this finding does align with our underlying assumption that the ability to infer candidates' private information about ability is a key part of the search process. But Vinkenburg et al.'s (2014) claim that "the process by which [top] managers are selected remains uncharted territory" appears accurate, and the small literature that does exist has not addressed confidentiality in the search process.[4]

The paper is organised as follows. In Section 1 we present the model with continuous abilities. In Section 2 we present the model with high-and low-ability agents. We derive and characterise the equilibrium allocations; we also look at the implied salary-productivity relationship, the planner's problem, and the implications of multiple firms. Finally, we present our conclusions in Section 3, and offer some suggestions for further research.

2 Model 1: Continuous Abilities

There is one firm and two job candidates, indexed by $n = 1, 2$. The candidates are endowed with a level of ability, a_n, drawn from a uniform distribution on $[1, \alpha]$. Ability levels are initially private information, meaning each candidate knows only their own ability and the firm does not know the ability of either candidate. The distribution of abilities is public information, known by both candidates and by the firm. The ability level of a candidate is revealed to the firm, before hiring takes place, if and only if that candidate applies and the hiring process is open.

Candidates are risk neutral and care about the value of the compensation they receive (which we refer to as the salary). They also care about the "stigma" that they would endure if they are

[4] Vinkenberg et al. provide a "conceptual" discussion of the selection of managers, focusing largely on internal selection, and do not consider confidentiality issues.

seen to apply for the position and be rejected. We model this as a utility cost $\sigma > 0$.[5] We assume that both candidates have the same outside option (that is, the value of their current employment situation), normalised to zero. The (value of the) output of the firm, if it successfully hires, is:

$$y = a_n;$$

if the firm does not hire, output equals zero. The model thus has only two parameters: α (the upper support of the ability distribution) and σ (the stigma of public rejection).

The game has three stages. In the first stage, the firm simultaneously announces both the type of the hiring process (open or secret search) and the level of the posted salary, which we denote by w. In the second stage, each candidate, given the realisation of their ability, chooses whether or not to apply (at which point they become an "applicant"). In the final stage, the firm chooses the applicant that maximises expected profit, given the information available to it; we refer to a hired applicant as an "employee".

2.1 Equilibrium

We restrict attention to symmetric subgame perfect Bayesian Nash equilibria. We use backwards induction, starting with the final stage of the game.

2.1.1 Stage 3: The Firm Chooses From the Applicant Pool

At this stage the firm takes as given its choice of hiring process, the posted salary, the candidates' application decisions, and the information available to the firm. If there are no applicants, then the vacancy is unfilled, and the payoff to the firm is zero. If only one candidate applies the firm hires that applicant at the posted salary level. If both candidates apply and the firm is unable to discern their abilities (because it adopted a secret search process) then the firm randomly chooses one of the two, with each candidate being chosen with probability 1/2. If both candidates apply, the firm

[5] We consider σ to be a placeholder for a variety of costs associated with being seen to apply for a job and being rejected. These include reputational costs, but also the possibility of negative repercussions at a candidate's current place of employment (if, for example, the candidate is deemed to be "disloyal"), and so on. It is also worth noting at this point that both levels of ability are "suitable for hire", or "appointable", in the sense that they provide more surplus than the firm's alternative payoff when it is unable to hire at all (which is set to zero). If we introduced unappointable candidates (that is, candidates with negative payoffs) into this analysis, secret search would be less attractive for the firm. With open search, however, unappointable candidates would never apply in equilibrium – because they would be rejected, and pay the penalty σ with certainty, if they did.

chooses the applicant with the higher expected payoff.[6]

2.1.2 Stage 2: Candidates Choose Whether or Not to Apply

At this stage, taking as given the firm's hiring process and the posted salary, candidates simultaneously choose their probability of applying. We consider secret search and open search in turn.

Candidates' Decisions Under Secret Search

Candidates' decisions under secret search are straightforward. Under secret search, applicants' identities are concealed and so neither candidate bears the risk of paying the cost σ. Since each candidate's outside option is normalized to zero, there are no other costs of applying, and there is a non-zero probability of being offered the job, the dominant strategy for a candidate—irrespective of ability—is to apply with probability one for any $w > 0$. For simplicity of exposition we assume that agents will also apply when they are indifferent ($w = 0$).

Candidates' Decisions Under Open Search

Consider the decision of agent 1. This candidate observes her own ability, a_1. It is immediate that the optimal strategy takes the form of a cutoff (apply whenever $a_1 > \bar{a}_1$), since if it is beneficial for her to apply for a given realisation of ability, it must also be beneficial to apply for all higher realisations. Agent 1 knows the salary and has a conjecture about the cutoff choice of candidate 2, $\bar{a}_2$. Given $\bar{a}_2$, we know that the probability that 2 applies is given by $\hat{z}_2 = (\alpha - \bar{a}_2)/(\alpha - 1)$ and the probability he does not apply is given by $1 - \hat{z}_2 = (\bar{a}_2 - 1)/(\alpha - 1)$. We determine agent 1's optimal cutoff by finding the value of a_1 for which her expected benefit from applying equals zero.

First, consider the case where agent 1's realised ability is below $\bar{a}_2$. Agent 1 gets the job only if agent 2 does not apply, which occurs with probability $1 - \hat{z}_2$. Her expected payoff is given by

$$\begin{aligned}(1 - \hat{z}_2)w + \hat{z}_2(-\sigma) &= \left(\frac{\bar{a}_2 - 1}{\alpha - 1}\right) w + \left(\frac{\alpha - \bar{a}_2}{\alpha - 1}\right)(-\sigma) \\ &= \left(\frac{w + \sigma}{\alpha - 1}\right)(\bar{a}_2 - \tilde{a}),\end{aligned}$$

where $\tilde{a} = (w+\alpha\sigma)/(w+\sigma)$. It follows that 1 should apply if $\bar{a}_2 > \tilde{a}$. Next, suppose agent 1's realised ability is above $\bar{a}_2$. In this case, if she applies, she gets the job with probability $(a_1 - 1)/(\alpha - 1)$

[6] We assume at this stage that the posted salary is not high enough to make expected profits negative *ex post*, and we verify later that this assumption is always justified in equilibrium.

(which is the probability that 2's realisation is below a_1), and her expected payoff is

$$\left(\frac{a_1 - 1}{\alpha - 1}\right) w + \left(\frac{\alpha - a_1}{\alpha - 1}\right)(-\sigma) = \left(\frac{w + \sigma}{\alpha - 1}\right)(a_1 - \tilde{a}).$$

In this case, agent 1 should apply provided that $a_1 > \tilde{a}$.

Putting these together, it follows that agent 1 should apply unless $a_1 < \tilde{a}$ *and* $\bar{a}_2 < \tilde{a}$. Translating this into a cutoff strategy, we obtain:

$$\bar{a}_1 = \tilde{a} \ \forall \bar{a}_2 \leq \tilde{a}$$

$$\bar{a}_1 = 0 \ \forall \bar{a}_2 > \tilde{a}.$$

The cutoff is decreasing in the salary: a higher salary leads to a greater likelihood of applying. The cutoff is increasing in the stigma: a greater disutility of rejection reduces the likelihood of applying. The symmetric Nash equilibrium of this game occurs where $\bar{a}_1 = \bar{a}_2 = \tilde{a}$. In Nash equilibrium, the probability that an agent applies is given by

$$z = \frac{\alpha - \tilde{a}}{\alpha - 1} = \frac{w}{w + \sigma}.$$

2.1.3 Stage 1: The Firm Chooses the Process and Salary

We now find the profit functions for the firm under both open and secret search processes.

Secret Search

With a secret search strategy, the firm optimally sets the salary to zero. Both candidates apply, but the firm cannot observe abilities. It randomly hires one applicant, with expected profit given by

$$\Pi^{ss} = \frac{1 + \alpha}{2}.$$

Open Search

With an open search strategy, the firm chooses a salary to influence the probability that candidates will apply for the job. With probability $(1 - z)^2$, there are no applicants; with probability $2z(1 - z)$, there is one applicant; and with probability z^2, there are two applicants. The firm's expected profit is given by

$$\Pi^{os} = 2z(1-z)\left(\frac{\alpha+\tilde{a}}{2} - w\right) + z^2\left(\frac{2\alpha+\tilde{a}}{3} - w\right)$$

where the second term uses the expected maximum value of two independent draws from the uniform distribution. Substituting the solution for z, we obtain

$$\Pi^{os} = \frac{2w\sigma}{(w+\sigma)^2}\left(\frac{\alpha+\tilde{a}}{2} - w\right) + \frac{w^2}{(w+\sigma)^2}\left(\frac{2\alpha+\tilde{a}}{3} - w\right),$$

which, after some algebra, becomes:

$$\Pi^{os} = \left(\frac{w}{3(w+\sigma)^3}\right)\left[6\sigma\alpha w + 6\sigma^2\alpha - 9\sigma w^2 - 6\sigma^2 w + 3\sigma w + 2\alpha w^2 - 3w^3 + w^2\right].$$

The firm chooses w to maximise this (rather complicated) expression for expected profit. In the Appendix we give a sufficient condition for this function to be strictly concave, and indeed simulations confirm that it has a unique maximum in $[0, \tilde{a}]$.

2.2 Implications

Figure 1 illustrates the locus $\Pi^{ss} = \Pi^{os}$ that divides the σ-α space into the regions where the firm chooses secret search and open search. As intuition would suggest, open search is preferred (a) for high α—because information about ability is more valuable to the firm—and (b) for low σ—because a high disutility from rejection implies the firm must pay a higher salary. Consistent with this logic, simulations reveal that the optimal salary under open search is increasing in α and in σ.

Other things equal, we know that a higher salary makes candidates more likely to apply ($\tilde{a}$ is decreasing in w) An implication is that a higher salary reduces the expected quality of applicants: lower quality candidates are more willing to apply when the salary is higher. However, the firm's expected quality of *employee* goes up, because of the increased probability that the firm will be able to select from more than one applicant.

2.3 The Planner's Problem

The equilibrium allocation is not always socially optimal. To show this, we consider a constrained planner's problem where the planner cannot observe abilities but can dictate the hiring strategy of the firm. That is, the planner chooses (a) whether the search process is confidential and (b) the posted salary. This is a policy that, in principle at least, would be implementable in the real world.

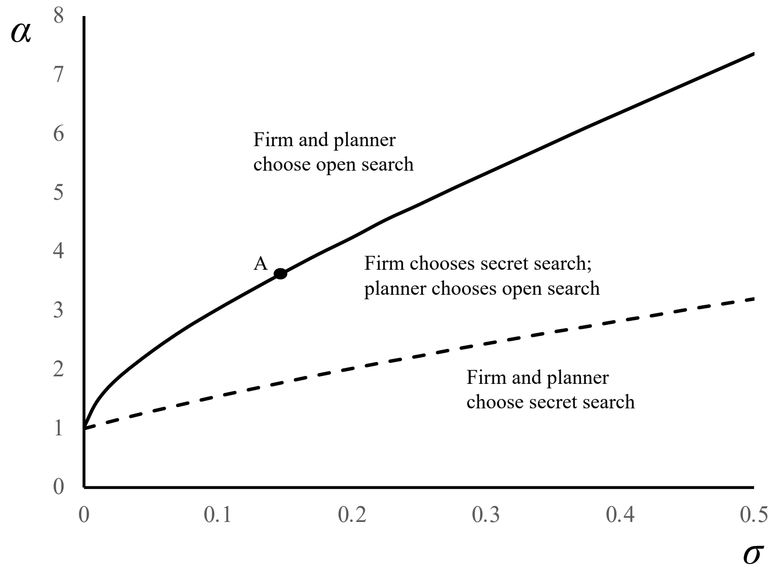

Figure 1: The firm's choice of strategy in σ-α space.

Thus, consider the problem of a planner who chooses the regime (secret search or open search) and the optimal salary, in order to maximise expected output minus the expected stigma of applicants, taking as given the response of candidates.[7]

Under secret search, the outcome of the planner's problem is the same as in the decentralised problem: both agents apply and there is no stigma from rejection, so welfare simply equals expected output from a randomly chosen agent ($W = \Pi^{ss} = (1+\alpha)/2$). The choice of salary is irrelevant here; a non-zero salary would simply be a redistribution from the firm to the employee, with no efficiency implications.

Under open search, the firm maximises expected output minus expected cost. It hires when at least one agent applies, so

$$\text{Expected cost (firm)} = [1-(1-z)^2]w = zw + z^2\sigma$$

[7]There are of course other versions of the planner's problem where the planner is less restricted. It would be interesting to consider whether there are mechanisms available to a planner that might eliminate the inefficiency that arises when there are no applicants, for example.

(using the solution $z = w/(w + \sigma)$ and rearranging). If we assume that the planner has access to the same technology as the firm—that is, they must set a salary to induce agents to apply, but can then observe ability, the planner would maximise expected output minus expected stigma. From the planner's point of view, the only cost arises if both agents apply, in which case one agent will be rejected and suffer the stigma. The cost from the planner's point of view is therefore

$$\text{Expected cost (planner) } = z^2\sigma$$

Comparing the two expressions, we see that the firm, like the planner, compensates for expected stigma, but also bears an additional cost equal to zw. Given that zw is strictly increasing in w, it follows that, at the private optimum, the planner's objective function has positive slope: the planner would strictly prefer a higher salary in order to make it more likely that at least one agent will apply. While the firm sees the salary as a cost, the planner views a higher salary simply as a transfer from the firm to the worker, with no efficiency implications.

A corollary is that there are regions of σ-α space where the firm would choose secret search, but where open search is socially preferred. The intuition is most easily understood from an envelope theorem argument. Imagine that the firm is at point A in Figure 1, where it is indifferent between secret search and open search. Now, starting from the privately-optimal open search equilibrium, consider the effect of a marginal increase in the salary. By the envelope theorem, this has only a second-order effect on the firm's profits. However, we have just shown that this change would strictly increase welfare, so the planner strictly prefers open search at this point. Figure 1 shows the boundary between the regions where the planner would choose secret versus open search.

The model thus has two main welfare implications. First, there are circumstances in which employees are *underpaid* relative to the social optimum. If we think of this as a model that might be applicable to the search for CEOs, this is perhaps a surprising conclusion. But it makes perfect sense in the context of the model: the firm has to bear the cost of encouraging candidates to apply, and does not, at the margin, receive all the associated benefit. The second implication has to do with how this underpayment can also affect the choice of regime. For some parameter values, firms will adopt either secret search even though open search is socially preferred. A high salary allows firms to identify high-quality applicants from a talent pool that is larger (in expectation). Thus, firms are both relatively likely to attract at least one applicant and can identify the higher-quality applicant if both candidates apply. The *ex post* suboptimal outcome of failing to hire is less likely to occur. But open search is costly, because of the high salary involved, so the firm may instead choose a confidential hiring process, taking the chance of either hiring a lower-ability candidate or even failing to hire at all, in order to be able to pay a lower salary, even though this is socially suboptimal.

2.4 Multiple Applicants

Finally, we consider what would happen if there were more than two candidates. Suppose there are $N+1$ candidates, and consider the symmetric case where all but agent i have chosen a cut off equal to $\bar{a}$. The probability of applying for each of these agents, as before, is given by $\hat{z} = (\alpha - \bar{a})/(\alpha - 1)$.

As before, suppose first that agent 1 observes a realisation of her own ability equal to a_1. In the case where $a_1 < \bar{a}$, she would get the job only if no one else applies, which occurs with probability $(1-\hat{z})^N$. The expected payoff to the applicant is given by

$$(1-\hat{z})^N w + [1-(1-\hat{z})^N](-\sigma) = (1-\hat{z})^N (w+\sigma) - \sigma$$

which is positive if

$$\begin{aligned}\left(\frac{\bar{a}-1}{\alpha-1}\right)^N &> \frac{\sigma}{w+\sigma}\\ \Rightarrow \bar{a} > \tilde{a} &= 1 + \left(\frac{\sigma}{w+\sigma}\right)^{\frac{1}{N}} (\alpha - 1)\\ \Rightarrow \hat{z} &= 1 - \left(\frac{\sigma}{w+\sigma}\right)^{\frac{1}{N}}.\end{aligned}$$

(It is easily confirmed that this corresponds to our earlier expression for the case of two agents when $N=1$.) If agent 1 observes $a_1 > \bar{a}$, she will receive the job if $a_1 > a_j \ \forall j \neq 1$, which occurs with probability $[(a_1-1)/(\alpha-1)]^{1/N}$. Similar reasoning to before yields that agent i should apply if $a_i > \tilde{a}$. The cutoff, $\tilde{a}$, is increasing in N: with more agents, it is less likely that any individual will want to apply at a given salary.

The firm's expected profit is given by

$$\Pi^{os} = \sum_{k=1}^{N+1} \binom{N+1}{k} z^k (1-z)^{N+1-k} \left[\left(\frac{1}{k+1}\right)\tilde{a} + \left(\frac{k}{k+1}\right)\alpha - w\right].$$

This is a standard binomial expansion, but with the first term (which would usually have coefficient $(1-z)^{N+1}$) omitted because this corresponds to the case where all candidates receive a draw below $\bar{a}$ and so the firm receives no applicants. We noted that, holding the salary fixed, an increase in the number of candidates reduces the probability that any individual agent will apply. Simulations reveal that this dominates the direct effect of an increase in N: the expected number of candidates goes down (still holding the salary fixed). The firm responds to this by increasing the salary, which increases the expected number of candidates and, despite the higher salary, expected profits increase. Of course, increasing the number of agents makes no difference under secret search: the firm still has expected profit equal to the expected output of a randomly selected agent ($\frac{1+\alpha}{2}$). Therefore the benefit from open search increases as we increase the number of agents: the loci in Figure 1 shift down.

3 Model 2: Discrete Abilities

The model of the previous section provides insight into when firms would choose secret or open search. However, it ignores a critical concern for both firms and candidates: the nature of the applicant pool. Firms, when choosing their search process, may think the applicant pool is relatively homogeneous, in which case they may see little benefit in open search; conversely, if the applicant pool is diverse, they may see great benefit in acquiring additional information. By the same token, candidates who are worried about the stigma of rejection need to think about the quality mix of the overall applicant pool, because this affects their likelihood of success if they apply. In this section, we consider a second version of our model in order to address such questions.

As before, we have one firm and two job candidates, indexed by $n = 1, 2$. In this variant, the candidates' ability takes one of two possible values:

$$a_n \in \{1, \alpha\} \quad \text{where } \alpha > 1.$$

We call candidates with $a_n = 1$ "low ability" and those with $a_n = \alpha$ "high ability". As in Model 1, ability levels are initially private information, meaning each candidate knows only their own ability and the firm does not know the ability of either candidate. The probability that either candidate is high-ability is public information, known by both candidates and by the firm. Specifically,

$$\begin{aligned} q &= \Pr(a_n = \alpha); \\ 1 - q &= \Pr(a_n = 1). \end{aligned}$$

The ability level of a candidate is revealed to the firm, before hiring takes place, if and only if that candidate applies and the hiring process is open. The other assumptions of the model are unchanged.

3.1 Equilibrium

The third stage of the game is unchanged, except that it is possible under open search that there will be two applicants with identical abilities. In this case, the firm randomly chooses one of the two.

3.1.1 Stage 2: Candidates Choose Whether or Not to Apply

At this stage, taking as given the firm's hiring process and the posted salary, candidates simultaneously choose their probability of applying. As in the previous version of the model, the dominant strategy for a candidate of either type under secret search is to apply with probability one for any $w \geq 0$. Under open search, however, the candidates' problem is more complex.

Candidates' Decisions Under Open Search

In general, candidates can choose mixed strategies. Let $z_H \in [0, 1]$ denote the strategy of a high-ability candidate: z_H is the probability that a high-ability agent chooses to apply for the position. Similarly, let $z_L \in [0, 1]$ denote the strategy of a low-ability candidate: z_L is the probability that a low-ability agent chooses to apply. The candidates take the value of the posted salary as given.

A candidate's decision about whether to apply depends on her probability of getting the job. This in turn depends on her own ability (which she knows) and the ability and strategy of the other agent (which she does not know). Consider a high-ability applicant. Define $\pi_H(\hat{z}_H, \hat{z}_L)$ as the probability that this applicant gets the job, conditional on the application strategy of the *other* agent—which in turn depends upon whether that other agent is high-ability or low-ability. That is, $\pi_H(\hat{z}_H, \hat{z}_L)$ gives the high-ability applicant's probability of being offered a job if the other agent is high ability and chooses $\hat{z}_H$ or if the other agent is low ability and chooses $\hat{z}_L$. This probability is defined conditional on the agent being an *actual* applicant for the job (that is, when the realisation of her mixed strategy results in her applying), but before it is known whether the other agent applies. Likewise, define $\pi_L(\hat{z}_H, \hat{z}_L)$ as the probability that a low-ability applicant gets the job, conditional on the possible application strategies of the other agent. The functions have the following properties.

1. $\pi_H(\hat{z}_H, \hat{z}_L)$ is decreasing in its first argument and independent of its second argument. If the other agent is also high-ability, then an increase in that agent's likelihood of applying (an increase in $\hat{z}_H$) makes it less likely that the applicant will get the job. If the other agent is

low-ability, then the high-ability applicant is guaranteed to get the job no matter what the value of $\hat{z}_L$. Both claims follow from the firm's equilibrium strategies in the third stage.[8]

2. $\pi_L(\hat{z}_H, \hat{z}_L)$ is decreasing in both arguments. A greater likelihood that the other agent applies, be she high- or low-ability, reduces the likelihood that a low-ability agent will get the job.

3. $\pi_L(\hat{z}_H, \hat{z}_L) < \pi_H(\hat{z}_H, \hat{z}_L) \forall \{\hat{z}_H, \hat{z}_L\} \in [0,1]^2 \backslash \{0,0\}$. That is, for any given strategy pair for the other agent, a high-ability applicant will always be at an advantage, *ex ante*, relative to a low-ability applicant. Again, this follows from the firm's strategies in the third stage: there are circumstances where a high-ability applicant is more likely to be hired than a low-ability applicant, but none where a high-ability applicant is *less* likely to be hired.[9]

4. $\pi_H(1,0) > \pi_L(1,0) > \pi_L(1,1)$. These inequalities follow from the above properties. Specifically, it is easy to show that $\pi_H(1,0) = (2-q)/2$, $\pi_L(1,0) = 1-q$, and $\pi_L(1,1) = (1-q)/2$.

Armed with these preliminaries, we now solve for optimal strategies of potential applicants, after which we identify equilibria in this subgame. A candidate's expected profit from applying, under open hiring, is $\pi_i()w - (1-\pi_i())\sigma$ for $i = H, L$. If an agent does not apply, she gets her reservation value of zero. Therefore, given that agents are risk-neutral, the optimal strategies of high-ability and low-ability agents are as follows:

$$If\, \pi_i(\hat{z}_H, \hat{z}_L)w - (1-\pi_i(\hat{z}_H, \hat{z}_L))\sigma \left\{ \begin{array}{ccc} >0 & then & z_i^* = 1 \\ =0 & then & z_i^* \in [0,1] \\ <0 & then & z_i^* = 0 \end{array} \right\} for\ i = H, L$$

where a star indicates an optimally-chosen strategy.

Lemma 1 *1. For a given pair $\{\hat{z}_H, \hat{z}_L\} \in [0,1]^2 \backslash \{0,0\}$, if high-ability candidates are optimally choosing not to apply or are optimally playing a mixed strategy, low-ability agents will choose not to apply ($z_H^* < 1 \Rightarrow z_L^* = 0$).*

2. For a given pair $\{\hat{z}_H, \hat{z}_L\} \in [0,1]^2 \backslash \{0,0\}$, if low-ability candidates are optimally playing a mixed strategy or are choosing to apply with probability one, high-ability candidates will choose to apply with probability one ($z_L^* > 0 \Rightarrow z_H^* = 1$).

Proofs for this Lemma and all subsequent results are in the Appendix.

[8] We could obviously specify a function with one argument only, but choose this form for clarity of exposition.

[9] There is no advantage to being high ability when $\hat{z}_L = \hat{z}_H = 0$; the applicant will be hired with certainty.

Proposition 1 *There is a unique Nash equilibrium in the second-stage game. The equilibrium strategies of applicants depend upon the value of the posted salary, as follows:*

1. *If* $0 < w < \sigma q/(2-q)$ *then low-ability candidates do not apply* ($z_L^* = 0$)*, and high-ability candidates apply with probability* $z_H^* = 2w/(q(w+\sigma)) \in (0,1)$.

2. *If* $\sigma q/(2-q) \leq w \leq \sigma q/(1-q)$ *then low-ability candidates do not apply* ($z_L^* = 0$)*, and high-ability candidates apply with probability one* ($z_H^* = 1$).

3. *If* $\sigma q/(1-q) < w < (\sigma(1+q))/(1-q)$ *then low-ability candidates apply with probability* $z_L^* = 2\left[1 - \sigma/((w+\sigma)(1-q))\right] \in (0,1)$*, and high-ability candidates apply with probability one* ($z_H^* = 1$).

4. *If* $(\sigma(1+q))/(1-q) \leq w$ *then both low-ability and high-ability agents apply with probability one* ($z_H^* = z_L^* = 1$).

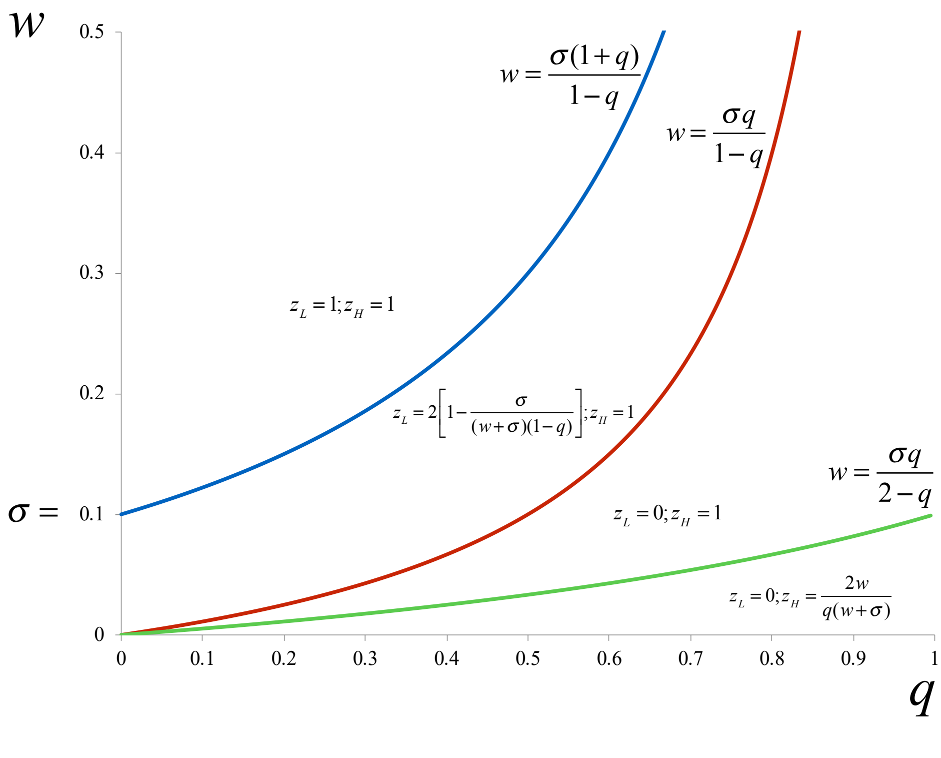

Figure 2: Equilibrium application strategies under open search

Figure 2 illustrates these decisions in q-w space (with $\sigma = 0.1$). As the average quality of the labour pool improves (q increases), both low-ability and high-ability candidates know that it is more likely that the *other* is high ability. Thus, other things equal, higher values of q deter candidates from applying. Increases in the salary make it more attractive for candidates to apply, while increases in σ have the opposite effect (in the figure, this would show up as upward movements of the boundaries). To summarise, the characteristics of the equilibrium of the second-stage game depend on the salary. By raising the salary, the firm can increase the probability that candidates apply. The strategic decision of the firm in the first stage of the game is then to decide whether or not to conduct an open search process and, if so, how much to pay to attract applicants of different types.

3.1.2 Stage 1: The Firm Chooses the Process and Salary

We now find the profit functions for the firm under both open and secret search processes. While secret search is straightforward, open search is more complicated. Proposition 1 reveals that if the firm elects to use open search, it has two possible options. It could offer a relatively low salary ($w \leq \sigma q/(1-q)$), for which low-ability candidates do not want to risk applying and being rejected. In this case, only high-ability candidates apply with positive probability. Alternatively, the firm could offer a relatively high salary ($w > \sigma q/(1-q)$), such that high-ability candidates apply with certainty, and low-ability candidates apply with positive probability. We thus have a formal definition of the two regimes that we introduced in our introduction: we refer to $w \leq \sigma q/(1-q)$ as *open separating* and we call $w > \sigma q/(1-q)$ *open pooling*. For expositional purposes, we discuss the two cases separately.[10]

Secret Search

Under secret search candidates apply with probability one if they obtain their reservation salary, so the firm chooses $w = w_{ss}^* = 0$. Because both agents apply with certainty, there are three possible outcomes of the second-stage game: two high-ability applicants (HH), two low-ability applicants (LL), or one applicant of each type (HL). In the HH case the firm hires a high-ability agent with certainty; in the LL case the firm hires a low-ability agent with certainty. If there is one of each, the firm cannot discern who is better and so hires the high-ability agent with probability 1/2.

The firm's profits are given by $y_n = a_n - w_{ss} = a_n$. Conditional on the realisation of applicants'

[10] In what follows, we refer to the choice of regime as a strategic decision of the firm, although in the end the firm simply chooses the generically unique salary (and implied regime) that maximises expected profit. Note also that the term "pooling" here carries a slightly different sense to the standard imperfect information story: even in our pooling regime, the firm can, by the assumption of open search, still distinguish between the two types of candidate.

abilities, this gives the following expressions for profit:

$$\begin{aligned}\Pi_{HH} &= \alpha \\ \Pi_{HL} &= \frac{1+\alpha}{2} \\ \Pi_{LL} &= 1\end{aligned}$$

The first outcome occurs with probability q^2, the second with probability $2q(1-q)$, and the third with probability $(1-q)^2$. Thus the firm's expected profit from secret search is

$$\Pi^{ss} = q\alpha + (1-q)$$

which simply equals the expected ability of a randomly chosen applicant.

Open Separating Strategy

If the firm wishes to use an open separating strategy, it chooses $w = w_{os} \leq \sigma q/(1-q)$. Given that search is open, the advantage of this separating strategy is that the firm can pay a relatively low salary. The disadvantage, relative to a pooling strategy, is that it will not attract low-ability applicants, which means that it runs a higher risk of having no applicants at all.

There are three possible outcomes of the second-stage game: two high-ability applicants (HH), one high-ability applicant (Hϕ), or no applicants ($\phi\phi$). Let w_{os} be the salary the firm posts under an open separating strategy. The firm's profits, conditional on the realisation of applicants' actual abilities and application choices, are given by

$$\begin{aligned}\Pi_{HH} &= \alpha - w_{os} \\ \Pi_{H\phi} &= \alpha - w_{os} \\ \Pi_{\phi\phi} &= 0.\end{aligned}$$

The probability of attracting no applicants equals $(1-qz_H)^2$. Define the probability of attracting at least one high-ability applicant as

$$\psi(w) = 1 - (1-qz_H)^2 = qz_H(2-qz_H).$$

where we make explicit the fact that $\psi()$ depends on the salary (because a high-ability agent's choice of strategy in the second stage game (z_H) depends upon the salary); specifically, $z_H = min\,(1, 2w/(q(w+\sigma)))$.

Using this definition, expected profit under the open separating strategy is

$$\Pi^{os}(w_{os}) = \psi(w_{os})(\alpha - w_{os})$$

Lemma 2 *Let* $w^*_{os} = argmax\,\Pi^{os}(w_{os})$*. If* $\alpha < \sigma q/(1-q)$*, the firm sets* $w_{os} = w^*_{os} = \sigma\alpha/(2\sigma+\alpha)$ *and high-ability applicants adopt a mixed strategy. If* $\alpha \geq \sigma q/(1-q)$*, then, under a separating strategy, the firm sets* $w_{os} = w^*_{os} = \sigma q/(2-q)$ *and high-ability applicants apply with probability one.*

Lemma 2 tells us that, depending on the values of the three parameters in the model, the firm may wish to induce high-ability candidates to apply with certainty, or only with positive probability. We can substitute in these values for the salary and calculate the value of expected profit at both an interior (mixed strategy) and corner (pure strategy) solution.

$$\text{For } \alpha < \frac{\sigma q}{1-q}, \Pi^{os} = \frac{\alpha^2}{\sigma+\alpha};$$
$$\text{For } \alpha \geq \frac{\sigma q}{1-q}, \Pi^{os} = \bar{\psi}[\alpha - w_{os}] = \bar{\psi}\alpha - \sigma q^2$$

where $\bar{\psi} = q(2-q)$ is the maximum value of $\psi(w)$; that it, it is the probability of there being at least one high-ability applicant when such applicants apply with probability one. Note that when high-ability agents adopt a mixed strategy, expected profit is independent of q. This is because the firm optimally chooses a value of the salary that is independent of q, while the optimal application strategy implies that the joint probability that an agent is both high-ability and chooses to apply ($qz_H = 2w/(w+\sigma)$) is likewise independent of q.

Open Pooling Strategy

Under open pooling, the firm chooses $w = w_{op} > \sigma q/(1-q)$. In this case, the firm sets a higher salary in order to induce low-ability applicants to apply as well. The benefit to the firm of such a strategy is that, if there are no high-ability candidates in the pool, it may still be able to attract a low-ability candidate, which is better than not hiring at all.

In this regime, high-ability agents apply with probability one, and low-ability agents apply with probability z_L, so all six possible combinations of applicants (HH, HL, LL, Hϕ, Lϕ, $\phi\phi$) are possible.

The firm's profits, conditional on the realisation of applicants' abilities, are

$$\begin{aligned}
\Pi_{HH} &= \alpha - w_{op} \\
\Pi_{HL} &= \alpha - w_{op} \\
\Pi_{LL} &= 1 - w_{op} \\
\Pi_{H\phi} &= \alpha - w_{op} \\
\Pi_{L\phi} &= 1 - w_{op} \\
\Pi_{\phi\phi} &= 0
\end{aligned}$$

The probability of getting at least one high-ability applicant is

$$\bar{\psi} = 1 - (1-q)^2 = q(2-q).$$

where we have imposed the condition that high-ability candidates apply with certainty ($z_H = 1$). Now define the probability of getting at least one low-ability applicant, *conditional on not attracting a high-ability agent*, as

$$\mu(w) = 1 - (1 - z_L)^2 = z_L(2 - z_L).$$

Expected profit under the open pooling strategy is then

$$\Pi^{op} = \bar{\psi}[\alpha - w_{op}] + (1-\bar{\psi})\mu(w_{op})[1 - w_{op}].$$

This profit function is more complicated than the one in the separating equilibrium and we no longer obtain a simple closed-form expression for the salary. The optimal open pooling salary at an interior equilibrium, w^*_{op}, is implicitly defined by

$$\frac{\partial \Pi^{op}}{\partial w_{op}} = (1-\bar{\psi})(1-w^*_{op})\mu'(w^*_{op}) - [\bar{\psi} + (1-\bar{\psi})\mu(w^*_{op})] = 0$$

We show in the Appendix (see the proofs of the comparative static results) that the second derivative of the expected profit function is everywhere negative, so it possesses a unique maximum. Further, we show that if the firm chooses to adopt a pooling strategy, it will always choose an interior equilibrium; that is, it will never set the salary so high that all candidates apply with probability one. Finally, we prove that $\mu'(w_{op})$ is positive, from which it follows that the equilibrium salary is less than one. This confirms that the firm always wants to hire *ex post*, even if it only has the option of hiring a low-ability applicant.

The Firm's Choice of Strategy

We can now characterise the firm's optimal choice in the first stage of the game. The profit functions that we have derived under the different hiring strategies are, in general, functions of the salary. As well as the direct dependence coming from the fact that the salary is a cost to the firm, there is an indirect dependence because the candidates' application strategies are also functions of the salary.

The firm can set a zero salary, and adopt secret search. Alternatively, it can select open search with a positive salary. As the firm varies the salary under open hiring, its expected profit function will generically have two local maxima, corresponding to the open pooling and open separating strategies. In stage one of the game the firm therefore chooses from {Secret search with $w = w^*_{ss} = 0$, Open search with $w = w^*_{os}$, Open search with $w = w^*_{op}$} in order to maximise expected profit.

3.2 Optimal Strategy of the Firm

We construct diagrams (see Figure 3) in q-α space that show the combinations of parameters where the firm adopts secret search, open search with a pooling strategy, and open search with a separating strategy. We do so by finding the loci where (a) the firm is indifferent between a pooling strategy and a separating strategy, (b) the firm is indifferent between secret search and a separating strategy, and (c) the firm is indifferent between secret search and a pooling strategy.

The locus defined by $\Pi^{op} = \Pi^{os}$ is a fixed value of q. To see why, note that

$$\bar{\psi}[\alpha - w^*_{op}] + (1-\bar{\psi})\mu(w^*_{op})[1 - w^*_{op}] = \bar{\psi}\alpha - \sigma q^2$$
$$\Rightarrow -\bar{\psi}w^*_{op} + (1-\bar{\psi})\mu(w^*_{op})[1 - w^*_{op}] = -\sigma q^2$$

We know that if $\alpha > \sigma q/(1 = q)$, w^*_{op} is independent of α, in which case the boundary is also independent of α. We confirm below (see Lemma 3) that this condition holds. At the boundary, the salary in the separating equilibrium is set sufficiently high to ensure that high-ability agents apply with probability one, so the expected profit from hiring high-ability agents is the same under either strategy—hence α is irrelevant.[11]

[11] The boundary does depend on the value of σ. Figure 3 is drawn for $\sigma = 0.1$, and the critical value is $q \approx 0.55$. We already know that an increase in σ leads firms to increase the salary in order to partially offset the negative effect on application probabilities. Under pooling, both high and low ability candidates require compensation, and low-

The other two loci are:

$$\Pi^{ss} = \Pi^{os} \Leftrightarrow \alpha = \frac{1}{q} + \frac{\sigma q}{1-q}$$
$$\Pi^{ss} = \Pi^{op} \Leftrightarrow \alpha = \frac{1}{q} + \frac{(2-q)w^*_{op}}{1-q} - \frac{1-q}{q}\mu(w^*_{op})[1-w^*_{op}].$$

The first expression is first decreasing, then increasing in q, attaining a minimum at $(1-\sigma^{1/2})/(1-\sigma)$. Because w^*_{op} depends in a complicated way on σ and q we cannot say more about the second locus analytically, but numerical simulations reveal that it is likewise initially decreasing and then increasing in q. The two loci intersect at the critical q that separates the pooling strategy from the separating strategy.

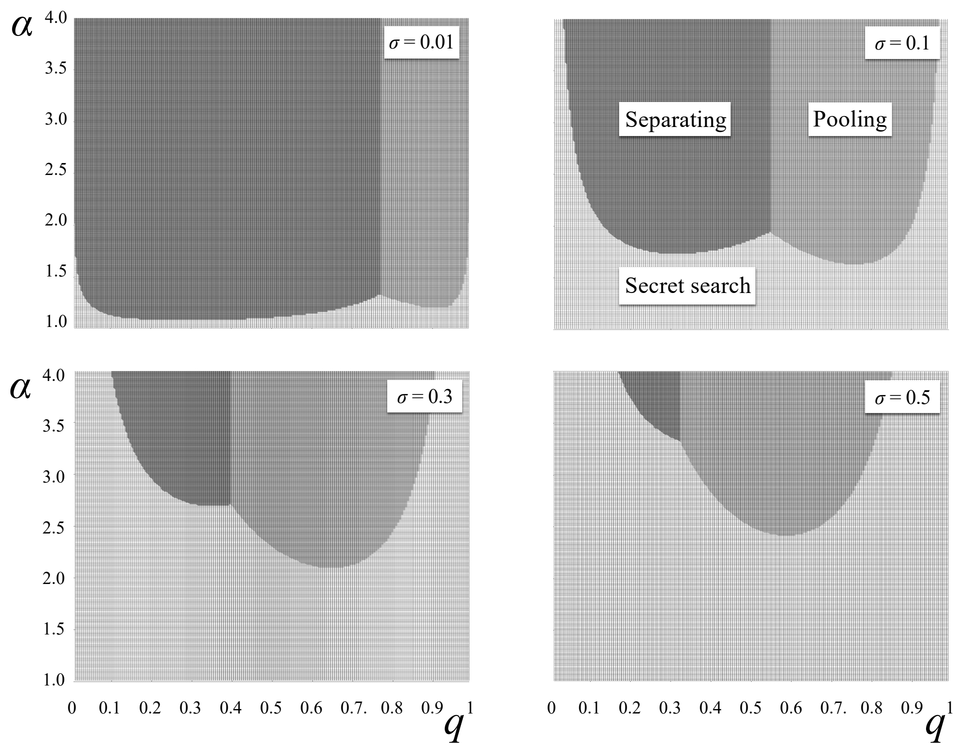

Figure 3: The firm's choice of strategy in q-α space.

The figure shows the firm's choice of strategy in a series of diagrams, with α and q on the axes,

ability candidates require more compensation, because they are more likely to incur the cost σ. Thus, as σ increases, the separating strategy (where only high ability candidates apply) becomes relatively more attractive to the firm. Numerical simulations confirm that the boundary does indeed move to the left as σ increases.

and σ as a shift variable.[12] As indicated in the top-right panel, the light, medium, and dark grey regions correspond to parameter values where secret search, open pooling search, and open separating search are optimal, respectively. We consider the full range of $q \in [0, 1]$, values of $\alpha \in (1, 4]$, and four different values of σ: $\sigma \in \{0.01,\ 0.1,\ 0.3,\ 0.5\}$. Because there are only three parameters in the model—α, q, and σ—we end up with a comprehensive characterisation of where each type of equilibrium obtains.

As one would expect, larger values of σ increase the area (marked in blue) in q-α space where secret search is preferred. The larger is σ, in general, the more compensation is required to induce candidates to apply under open search of either kind. Secret search requires no such compensation. Thus, with large values of σ, secret search is preferred by the firm provided that α is not too high. This is the same result as we found in the first version of the model.

For intermediate values of q, open search is preferred to secret search when α is large. Again, this confirms a finding from the previous model. Open search allows the identification of high-ability candidates, and the gain from detecting such agents is greater when α is large. However, at extreme values of q (either low or high), the candidate pool is close to being homogeneous. In these cases, the firm perceives little gain from the ability to distinguish candidate types because both agents are likely to be of the same type in any case. Thus the firm sees minimal benefit to paying the salary premium, and secret search dominates.

We can summarise our findings as follows. First, when there is not much variation in ability—in the sense that high-ability candidates are not much more productive than low-ability candidates—the firm chooses secret search, particularly when the stigma of being an unsuccessful applicant is high. Second, when candidates are largely homogeneous (q close to 0 or 1), the firm again chooses secret search, since there is unlikely to be much value in being able to distinguish applicants' ability. Third, in a world of "superstars"—big differences in ability, but with high ability candidates being rare—the firm selects a pooling strategy; it wants to identify high-ability candidates if possible, but also wants to insure against the possibility of having no high-ability applicants. Fourth, if high-ability candidates are relatively easy to find, the firm adopts a separating strategy; it expects to have a high-ability applicant, so does not deem it necessary to entice low-ability candidates to apply.

Finally, we offer a slight reinterpretation of the model that allows us to think about productivity of the firm as well as candidates. Suppose that not all firms are capable of making equal use of high-ability candidates. That is, imagine that a firm possess assets that are complementary to the high

[12]These figures were generated using MATLAB.

abilities of applicants, so that α reflects both the ability of the applicant and these complementary assets of the firm. Under this interpretation, high-α firms are those where talent is highly valued and utilised, while low-α firms are those where candidates are effectively relatively homogeneous. If we make the additional reasonable assumption that high-α firms are then more productive in general, we conclude that secret search will be chosen by relatively mediocre organisations. The choice of open search, by contrast, would be a signal that the firm sees value in discriminating on merit.[13]

3.3 Comparative Statics

We now turn to the comparative statics of the model. As a preliminary, recall that we showed in Lemma 2 that when $\alpha < \sigma q/(1-q)$, high-ability agents use a mixed strategy in the separating equilibrium. We have just seen, however, that for high values of q, the firm will choose secret search rather than open search. It turns out we never observe the mixed strategy in equilibrium: Lemma 3 shows that if the profit function under a separating strategy attains its maximum on the interior of $(0, \sigma q/(2-q)]$, the firm can always earn higher profits by adopting secret search. This implies in particular that if the firm chooses the separating strategy, it will always set a salary high enough to ensure that all high-ability agents apply. The key to the proof is the fact that maximum profits at an interior solution, $\Pi^{os} = \alpha^2/(\sigma+\alpha)$, are independent of q.

Lemma 3 *If a firm chooses a separating strategy in equilibrium, then it must be at the corner solution with* $w_{os} = w^*_{os} = \sigma q/(2-q) \Rightarrow z_H = 1$.

This allows us to ignore, in everything that follows, the cases where high-ability agents choose a mixed strategy. The proofs of all the comparative results in this section are in the Appendix.

3.3.1 Comparative Statics: Salary

We first consider how the salary is affected by changes in the parameters of the model. First, note that the optimal salary is independent of α. Under secret search, we know that the optimal salary is zero. And under pooling and separating, the firm sets a salary that ensures that high-ability firms apply with probability one—and so the salary is unresponsive to α at the margin. The salary is increasing in σ under open search, because higher levels of stigma must be offset by a higher salary.

The most interesting question is: how do changes in q affect the salary? First, we note that, in equilibrium, higher value of q are associated with higher salaries under both pooling and separating.

[13] Note that α does not need to be observable to the candidates *ex ante*; they simply respond to the posted salary.

The latter is immediate from the solution for w^*_{os}; the result for w^*_{op} is proved in the Appendix.

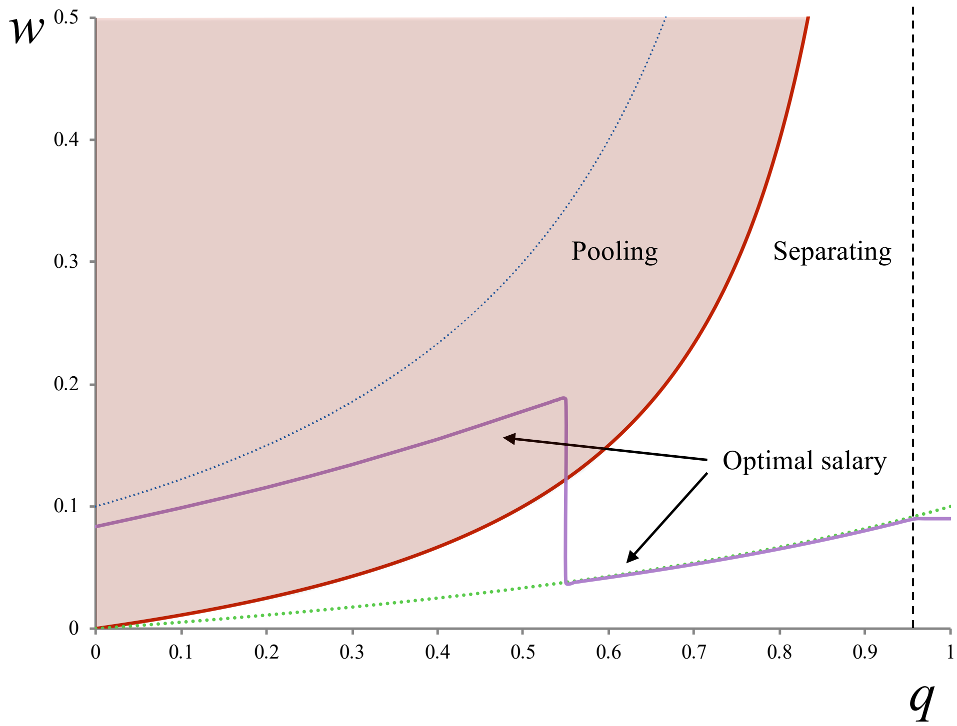

Figure 4: The optimal salary under open search as a function of q ($\sigma = 0.1$ and $\alpha = 1.8$)

Figure 4 builds on Figure 2 to show the optimal choice of salary under open search as q varies. At low values of q, the firm chooses a pooling strategy, because the risk of there being no high-ability agents in the mix is large. Note that the salary lies in the region where low-ability agents play a mixed strategy; as we have already observed, the firm will never find it optimal to set the salary so high that low-ability agents apply with probability one. As q increases, the firm sets a higher salary.

As we saw in Figure 3, there is a critical value of q at which the firm optimally switches to a separating strategy. At this point, the firm *reduces* the salary to the level that just guarantees that high-ability agents apply with certainty. Low-ability agents no longer find it worthwhile to apply. The relatively large value of q means that even with a separating strategy the firm is unlikely to end up with no applicants at all (recall that the probability of both candidates being low-ability is $(1-q)^2$) and so the firm finds it worthwhile to pay a lower salary and take this risk. Increases in q beyond this critical value again result in an increase in the salary, because this is needed to ensure that high-ability agents continue to apply with certainty. (There is a second critical value of q

above which the salary is constant, but—as already noted—this will not be observed in equilibrium because the firm will be using secret search.)

3.3.2 Comparative Statics: Expected Profits

Now consider expected profits. Not surprisingly, increases in α increase the firm's expected profits, while increases in σ reduce profits whenever the firm uses open search (because it has to pay a higher salary to overcome the greater stigma of being an unsuccessful candidate). The effect on expected profit of changes in q is less obvious. Holding constant the salary and candidates' application strategies, increases in q have the direct (partial) effect of increasing expected profit. But because increases in q make agents of both types more reluctant to apply, the total effect is unclear.

Under a separating strategy, profits are increasing in q.[14] In the open pooling regime, however, changes in q have an ambiguous effect on expected profit. In a pooling regime the firm may end up hiring either a high-ability candidate or a low-ability candidate. Increases in q shift the probability mass to hiring the former rather than the latter, thus increasing expected profit. This magnitude of this effect depends on how much more productive high-ability candidates are; specifically, it is proportional to $(\alpha - 1)$. But increases in q make it less likely that low-ability candidates will apply; this effect reduces expected profit and is independent of α. As a result, increases in q increase expected profit if α is sufficiently large, but decrease expected profit for small α.

3.3.3 Comparative Statics: Application Probabilities

Table 1 summarises the comparative statics results presented so far. It also includes the comparative statics for the application probabilities z_H and z_L, and—for comparison—the corresponding results from Model 1. In equilibrium, application probabilities are largely unaffected by changes in the parameters: high-ability agents always apply with probability one, while low ability agents apply with probability one under secret search and with probability zero in the separating equilibrium. In the pooling equilibrium, increases in either σ or q have the direct effect of reducing z_L. There is a partially offsetting effect from an increase in the salary, but the overall effect is still negative.[15]

[14] This follows immediately from substitution of the equilibrium value of the salary into the profit function. Note that, because we are at the corner solution with $z_H = 1$, the envelope theorem does not apply and changes in the salary have a direct effect on profit.

[15] We do not have analytical proofs of these two pooling results, but they are supported by extensive simulations.

Endogenous Variable	Derivative	Secret Search	Separating Equilibrium (when $\alpha \geq \sigma q/(1\text{-}q)$)	Pooling Equilibrium	Model 1 (Continuous Abilities)
Salary	$dw/d\alpha$	0	0	0	+
	$dw/d\sigma$	0	+	+	+
	dw/dq	0	+	+	n/a
Expected Profit	$d\Pi/d\alpha$	+	+	+	+
	$d\Pi/d\sigma$	0	–	–	–
	$d\Pi/dq$	+	+	+ or -	n/a
Application Probability	$dz_H/d\alpha$	0	0	0	+
	$dz_L/d\alpha$	0	0	0	
	$dz_H/d\sigma$	0	0	0	–
	$dz_L/d\sigma$	0	0	–	
	dz_H/dq	0	0	0	n/a
	dz_L/dq	0	0	–	

Table 1: Comparative Statics

3.3.4 Comparative Statics: Comparison Between the Models

In terms of comparative statics, the main difference between the two models is the following. In our model with continuous abilities, the application probability and the salary depended on α. More productive workers increase the cost of failing to hire, and increase the benefit of getting multiple applicants, so the firm responds to higher α by increasing the salary to encourage candidates to apply. But when agents are of discrete types, the firm always induces high-quality agents to apply with probability one, irrespective of the productivity gap between high- and low-quality candidates. Thus application probabilities and the salary are both unresponsive to changes in α.

3.4 The Salary-Productivity Relationship

In most economic models we expect to see a positive association between compensation and productivity. In the current setting, the relationship is more complicated and quite subtle. First, consider whether we expect to observe high-ability candidates earning more than low-ability candidates. Suppose we have observations from isolated labour markets, all with the same underlying parameters (q,

α, and σ). Then we would trivially always observe high-ability and low-ability candidates receiving the same salary, since—in any given regime—the firm offers the same salary to all applicants.

A second question is whether we expect to observe higher compensation in economies where α is higher. The answer, in the model, is yes and no. Higher α makes the zero-salary secret search regime less likely, just as we saw in Model 1. But the salary under open search is independent of α.

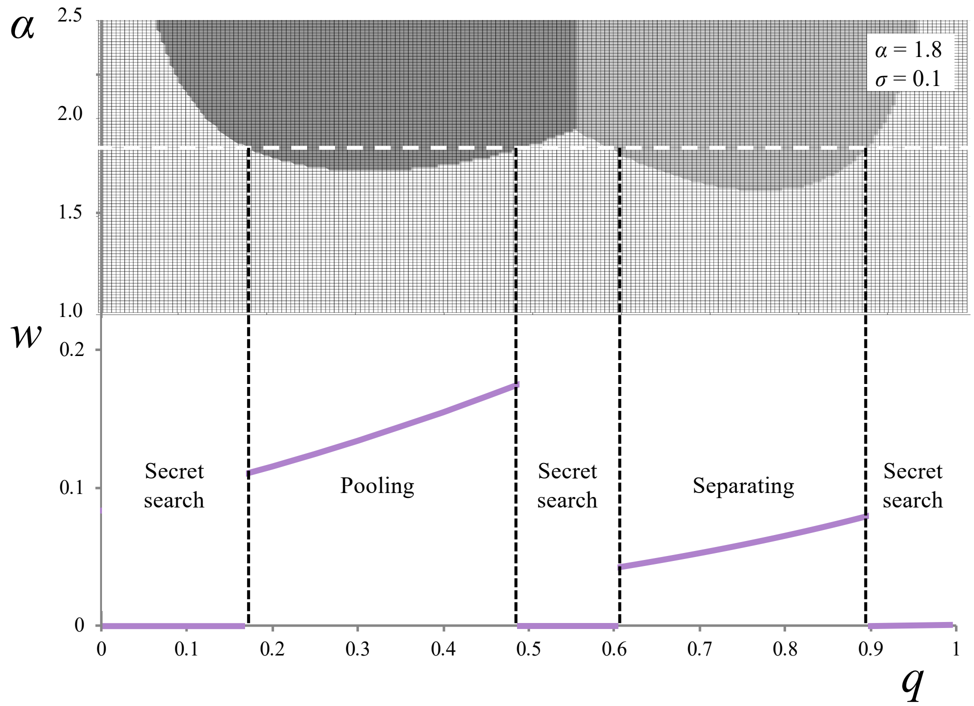

Figure 5: The effects of changes in the applicant mix (q) on the firm's choice of strategy and salary

A third question is whether we will observe higher salaries when the average quality of the labour pool increases—that is, when q is higher. Increases in q have the direct effect of improving the expected quality of the candidate pool. However, increases in q also make both types of candidates more reluctant to apply under open search, because each potential applicant recognises that there is a greater chance that the *other* candidate will be high-ability and will get the job. The firm therefore has to pay a higher salary in order to induce a given application strategy. The firm trades off all of these effects, with the result that there is a non-monotonic relationship between q and w and—for at least some values of α and σ—variations in q lead to multiple switches between regimes.

Figure 5 illustrates the choice of the firm as q varies when $\alpha = 1.8$ and $\sigma = 0.1$. The top panel shows a segment of the top right panel of Figure 3 while the bottom panel illustrates the salary paid by the firm (reproduced from Figure 4, and adjusted to include the zero salary paid under secret search). Tracing across from left to right, we start in the region where q is too small for open search to be worthwhile. The firm chooses secret search and the equilibrium salary is zero. As q increases, the fraction of high-quality candidates becomes large enough for the firm to switch to open search in order to distinguish quality types. However, at low q there is a significant risk that neither candidate is high-quality, so the firm chooses a high enough salary to induce the pooling equilibrium (indicated in red). There is a discrete jump in the salary, from zero to w^*_{op}, as q crosses this threshold. Further increments in q make both high- and low-ability candidates more reluctant to apply; the firm partially offsets this effect by offering a higher salary.

Once q reaches the second critical value in Figure 5, the salary necessary to sustain a pooling strategy becomes too high to justify the benefit of open search, yet the firm does not want to risk of attracting no candidates. Secret search is now the best option: even though the firm cannot identify a high-quality applicant, it knows that candidates will apply and it can pay a low (zero) salary. Further increments in q, however, bring the firm across the next threshold into the region where a separating strategy is the most profitable. The expected quality of the candidate pool is now so high that the firm is sufficiently confident of obtaining at least one applicant even when only high-quality candidates apply. Salaries jump up at the threshold from zero to w^*_{os}, which is positive, but lower than the minimal w^*_{op} in the pooling region, while further increases in q lead to higher salaries. When q crosses the final threshold, the likelihood that any candidate will be high quality is so high that the firm finds it more profitable to accept a random candidate through the secret search process, and salaries once again fall back to zero.

Within regions where the firm chooses secret search, there is no relationship between q and the salary, since the salary is zero for all values of q consistent with secret search. When changes in q induce shifts in the firm's choice of regime, we see higher q sometimes associated with higher salaries, and sometimes associated with lower salaries. But within both the pooling and separating regimes, there is a positive relationship between the average labour quality (q) and the salary.[16]

[16] A closely related question is: again allowing q to vary, what is the association between the salary and the expected productivity of a successful applicant? Under secret search, both the salary and expected ability are constant. In the separating regime, the salary changes as q changes, but any hired applicant is high-ability, so there is again no relationship. And in the pooling regime, there are two contradictory effects: as q increases, the randomly-drawn candidate is more likely to be high-ability, but a low-ability candidate is more likely to apply because of the higher salary. Simulations reveal that the first effect dominates: higher salaries are associated with higher expected productivity.

To summarise, our model at the individual level contains an efficiency-wage mechanism that is, to the best of our knowledge, novel in the literature. Firms offer a higher salary to attract a larger applicant pool—and they need to do so because of the cost that results from being an unsuccessful applicant. But, in the aggregate, this efficiency-wage effect does not necessarily translate into a positive relationship between salaries and productivity. Indeed, other things equal, higher salaries reduce the average quality of the applicant pool.

3.5 The Planner's Problem

As in our first model, the equilibrium allocation may be suboptimal. Under secret search, just as in the earlier model, there is no difference between the planner's solution and the decentralised solution. Under open search, the planner's problem involves maximising expected output minus the expected disutility of rejected agents, subject to the optimising behaviour of candidates:

$$\max_{w_{pp}} W = \psi(w_{pp})\alpha + (1-\psi(w_{pp}))\mu(w_{pp})$$
$$- \left[q^2(z_H(w_{pp}))^2 + 2q(1-q)z_H(w_{pp})z_L(w_{pp}) + (1-q)^2(z_L(w_{pp}))^2\right]\sigma$$

subject to

$$w_{pp} \geq 0$$
$$\psi(w_{pp}) = qz_H(w_{pp})(2 - qz_H(w_{pp}))$$
$$\mu(w_{pp}) = z_L(w_{pp})(2 - z_L(w_{pp}))$$
$$z_H(w_{pp}) = min\left[1, \frac{2w_{pp}}{w_{pp}(w_{pp}+\sigma)}\right]$$
$$z_L(w_{pp}) = min\left[1, max\left[0, 2\left[1 - \frac{\sigma}{(w_{pp}+\sigma)(1-q)}\right]\right]\right]$$

Proposition 2 confirms that the planner will, under some circumstances, want to set a higher salary than the firm and will also sometimes choose a different regime to the firm.

Proposition 2

1. *If the planner chooses a separating strategy ($w_{pp} \leq \sigma/(1-q)$), and if $\alpha < \sigma q/(1-q)$, she chooses the same salary as the firm: $w^*_{pp} = w^*_{os} = \alpha\sigma/(\alpha + 2\sigma)$. The allocation is identical to the firm's problem.*

2. *If the planner chooses a separating strategy ($w_{pp} \leq \sigma/(1-q)$), and if $\alpha \geq \sigma q/(1-q)$, she chooses a salary that ensures that high agents apply with probability one: $w^*_{pp} \in [\sigma q/(1-q), \sigma q/(2-q)]$.*

The allocation is identical to the firm's problem, except that output may be redistributed from the firm to the employee.

3. *If the planner chooses a pooling strategy ($w_{pp} > \sigma/(1-q)$), she chooses a higher salary than the firm: $w^*_{pp} > w^*_{op}$. The allocation implies higher expected output than in the firm's problem because there is a greater probability that low-ability agents will apply.*

4. *There exist parameter values such that the planner chooses a pooling strategy, even though the privately optimal hiring strategy is either secret search or a separating strategy.*

The intuition is similar, but not identical, to that of Model 1. Starting from a privately optimal open pooling equilibrium, consider the effect of a marginal increase in the salary. By the envelope theorem, this has only a second-order effect on the firm's profits. It also does not affect the expected utility of low-ability candidates because, at the open pooling equilibrium, low-ability candidates are indifferent about whether or not to apply—and this remains true following a small change in the salary. High-ability candidates, however, are inframarginal, so an increase in the salary leads to a first-order increase in their welfare. A marginal increase in the salary implies a possibility of paying a larger salary to high-ability workers, which, from the firm's perspective, is a cost with no associated marginal benefit, whereas the planner views this as a pure transfer with no efficiency implications.

As in our previous model, the firm will for some parameter values adopt secret search when a higher-salary pooling strategy is socially preferred. For some parameter values the firm will adopt open search, but choose a separating regime when a pooling regime is socially preferred. Pooling allows firms to identify high-quality applicants from a relatively large talent pool. But the firm may instead choose to adopt a separating strategy, taking the chance of failing to hire at all, in order to be able to pay a lower salary, even though this is socially undesirable.

3.6 Multiple Firms

Suppose we allow other firms to compete with each other in the market for candidates, where candidates choose where to apply. The environment is then that of a directed search model where, effectively, the firms sell jobs to candidates who have private information about their type. Considerable progress has been made in the mechanism design literature considering similar environments with large numbers of agents on each side of the market. In that literature, the focus has been on whether, in order to maximise profits, a seller should use simple price-posting (potentially inducing ex ante *sorting* equilibria, where high types and low types approach different sellers with different prices) or auctions (where high and low types approach each seller, and the seller *screens* ex post).

Cai et al. (2017) demonstrate that this decision depends crucially on the type of meeting mechanism assumed. Separate submarkets for each type of buyer (common with price posting models) are efficient if and only if meetings are bilateral. On the other hand, pooling buyers into a single market (common with competing auction models) is efficient if and only if the meeting process satisfies "joint concavity". The commonly used urn-ball matching process (which emerges as an equilibrium in, for example, Julien et al. 2000) satisfies joint concavity. [17] None of these results applies directly to the problem at hand. We are concerned here not just with private information about buyers and pricing but also with the penalty that a buyer may incur if the buyer attempts to purchase and is unsuccessful, where this penalty depends on the mechanism announced by the seller. To date, no such mechanism has been considered in this literature. Complicating matters further, the model in this paper is designed to analyse small markets for high-profile positions, whereas the papers on mechanism design in search environments typically focus on large markets where some of the difficult strategic considerations between sellers vanish in the limit. Analysing this issue head-on, then, is well beyond the scope of the present paper.

However, using the Cai et al. (2017) distinctions, but preserving the key margins that we focus on here, we can say that a model of this kind would most naturally be one with urn-ball meetings. To the extent that the Cai et al. (2017) results carry over into this more general environment, they would imply the use of auctions with *ex post* screening. In our context, this would mean that firms would sell off high-profile positions to candidates through a bidding process—something we do not typically see. Clearly, more work needs to be done in this area.

4 Conclusions

Our aim in this study was to investigate a straightforward trade-off: acquiring information about job applicants may be at odds with confidentiality in the search process. Confidentiality makes it easier to attract applicants, but harder to learn their true ability. The model we wrote down to consider this trade-off is intentionally very simple, in both its variants. Yet it yields surprisingly complex implications, many of which are not at all obvious. The unexpected richness of the model

[17] Numerous examples of both types of meeting processes and pricing mechanisms abound. For examples of bilateral meetings and price posting, see the competitive search models, following Moen (1987) with symmetric information and Guerrieri et al. (2010) with asymmetric information. for examples of urn-ball multilateral matching see McAfee (1993), Peters and Severinov (1997), Albrecht et al. (2014), and Auster and Gottardi (2017). The survey by Wright et al. (2017) provides an excellent discussion.

is due to the fact that the cost of applying for a position is endogenous. In our setting, potential applicants must assess their risk of being publicly seen to have been rejected for a position. But this risk depends upon the abilities and choices of other potential applicants. So the decision to apply for a position requires an assessment of the existence and quality of others who are competing for that position. This feature surely captures a genuine concern felt by high-profile applicants to high-profile positions.

Firms, in turn, must decide how to influence candidates' choices through choosing open or secret search and posting an appropriate salary. This decision depends upon the characteristics of the labour pool—both the relative superiority of high-ability candidates and—in our second model—the quality mix. The firm trades off three potential negative impacts on its profits: the risk of not hiring anyone at all; the opportunity cost of possibly hiring a relatively low quality applicant when there is a better applicant available; and the direct salary cost that must be paid to which ever applicant they hire. The magnitude of these tradeoffs depends on the search strategies of the candidates. In our second model, the tradeoffs depend particularly on the fact that low-ability candidates require higher salaries to induce them to apply than do high-ability candidates, because low-ability candidates face a higher probability of rejection. Secret search implies the lowest salary of all: there is no possible stigma from rejection and so the salary is set equal to the outside options of the candidates (normalised to zero in this model).

In our introduction we posed several questions; we now summarise the answers suggested by our model. Many of these answers come down to comparative static exercises with respect to the three parameters in the model: the utility cost of (observed) rejection, the relative productivity of a high-ability candidate, and the fraction of candidates that are of high ability.

Our first question was: what features of the candidate pool would make a firm choose secret search, open pooling search, or open separating search? If the cost of rejection is large, a firm will tend to choose secret search. With open search, a large disutility from rejection implies that the firm will need to pay a high salary to induce any candidates to apply. This is clearly costly, and will push the firm towards secret search to avoid these costs. For somewhat less straightforward reasons, if a firm does choose open search, a higher cost of rejection induces the firm to choose the open separating search rather than open pooling search. Intuitively, firms must pay a higher salary under open pooling than under open separating (to encourage low-ability candidates to apply). As the cost of rejection increases, the required increase in the pooling salary is larger than the corresponding required increase in the separating salary—so the proportion of the parameter space that supports open pooling search diminishes, relative to open separating search.

Higher values of relative productivity imply less secret search in both variants of the model. When high-ability employees are much more productive than low-ability employees, there is a large incentive to gather information about applicant types. This requires open search. The choice between open separating search and open pooling search is independent of relative productivity, however. This is because the salary in both cases is at or above the level that ensures that high-ability candidates apply with probability one, which—combined with the fact that the firm will always choose a high-ability applicant if available—implies the same expected profit from the hiring of a high-ability applicant.

Secret search is also more likely when the population of candidates is largely homogeneous, because there is then little gain from opening up the search process to get information about candidates' abilities. When the applicant pool is more mixed, open search is more likely. In a world of "superstars", where good candidates are valuable but rare, firms engage in open hiring and offer a high salary to insure against the possibility that they will fail to attract a high-ability candidate. Adopting a pooling strategy requires a higher salary, and reduces the expected quality of applicants relative to a separating strategy, but mitigates the risk of no-one applying. If high-ability candidates are relatively common, by contrast, the firm adopts a separating strategy (with a lower salary).

The second question we asked was: what type of candidate would be most likely to apply under each type of search? Under open separating search, only high-ability candidates apply. Under open pooling search, all high-ability candidates apply and some low-ability candidates apply. The probability of applying, for a low-ability candidate, is increasing with the salary—but it is never an equilibrium, under open pooling, for low-ability candidates to apply with certainty. Only under secret search do all candidates apply with certainty. Our first model yields similar findings: there is an ability cutoff, and only candidates with realisations above that cutoff will choose to apply.

The relationship between salary and productivity levels (the subject of our third question) is complex. Higher relative productivity, and a large pool of high-ability candidates, both imply higher expected productivity for the firm. Within the secret search region, the salary is zero and independent of relative productivity. Within the open pooling search region, the salary (which is already high enough to induce high-ability candidates to apply) is also independent of relative productivity, since it is being chosen to induce applications from low-ability candidates. And within the open separating search region, however, the salary is likewise independent of relative productivity, since it is always chosen to ensure all high-ability candidates apply. By contrast, in the model with continuous abilities, higher productivity of potential applicants does induce the firm to set a higher salary. When higher values of productivity are due to a higher quality mix of applicants, there is a complicated and non-monotonic relationship between productivity and salary, as illustrated in Figure 5. Within both the open pooling and the open separating regions, increases in the relative frequency of high-ability candidates are associated with higher salaries, but when changes in the applicant mix result in a change in hiring strategy, the salary jumps discretely, either upwards or downwards.

Our fourth question asked how equilibrium allocations compare with those chosen by a planner. We found that, whenever the planner chooses open separating search, the allocation is identical to that chosen by the firm. However, when the planner chooses open pooling search (or open search in our first variant), she chooses a higher salary than is optimal for the firm. With open pooling, high-ability employees collect rents—the salary is higher than needed to induce them to apply. Since part of the inducement to encourage low-ability candidates to apply is siphoned off (in expectation) by high-quality candidates, the marginal benefit that the firm gets from an increase in the salary is less than the marginal cost, and so the firm pays less than the socially optimal amount. Thus, from the planner's point of view, with open pooling search, firms should pay higher salaries—in order to induce more applications from low-ability candidates! The existence of these rents also implies that firms may choose a secret search regime even when it would be more efficient, from a social perspective, to adopt open hiring.

The last question we posed was: what kind of firms would choose secret search? We suggested an answer to this based on a slight reinterpretation of our model, where productivity results both from the employee's ability and from complementary assets of the firm. Under this interpretation, the model suggests that the types of firm that choose secret search are those with low productivity, where talent is not highly valued. Secret search becomes a signal of mediocrity of the organisation, while open search indicates a firm that has the ability to effectively utilise talented individuals.

Finally, because we have developed a very stylised and focused model, there are many extensions that we could consider. We discussed one in some detail: when we increase the number of firms the model becomes one of directed search with private information and urn-ball meetings. If firms search simultaneously, the existing literature suggests that firms should auction off jobs, with ex post screening—something that is rarely seen. However, if, instead, searches such as this are rare enough that competition among firms is sequential (rather than simultaneous) then the results of this model are more likely to carry over to this richer environment. Formal analysis of this environment is beyond the scope of this paper, but may be worth pursuing in subsequent work.

5 Appendix

Properties of the Profit Function under Open Search in Model 1

As preliminaries, note that

$$\begin{aligned}
\frac{\partial z}{\partial w} &= \frac{\sigma}{(w+\sigma)^2} > 0; \\
\frac{\partial^2 z}{\partial w^2} &= -\frac{2\sigma}{(w+\sigma)^3} < 0; \\
\frac{\partial \tilde{a}}{\partial w} &= -\frac{\sigma(\alpha-1)}{(w+\sigma)^2} < 0; \\
\frac{\partial^2 \tilde{a}}{\partial w^2} &= \frac{2\sigma(\alpha-1)}{(w+\sigma)^3} > 0.
\end{aligned}$$

The firm's expected profit is given by

$$\begin{aligned}
\Pi^{os}(w) &= 2z(w)(1-z(w))\left(\frac{\alpha+\tilde{a}(w)-2w}{2}\right) + z(w)^2\left(\frac{2\alpha+\tilde{a}(w)-3w}{3}\right) \\
&= \frac{1}{3}\left[z(w)(1-z(w))(3\alpha+3\tilde{a}(w)-6w) + z(w)^2(2\alpha+\tilde{a}(w)-3w)\right].
\end{aligned}$$

Differentiating with respect to w

$$\frac{\partial(3\Pi^{os})}{\partial w} = [z(3-2z)]\frac{\partial \tilde{a}}{\partial w} + [\alpha(3-2z)+\tilde{a}(3-4z)-6w(1-z)]\frac{\partial z}{\partial w} - 3z(2-z).$$

It is easily confirmed that this derivative is positive when evaluated at $w = 0$ and negative as $w \to \infty$. Now define $\bar{w} = ((1-\sigma)+[(1-\sigma)^2+4\alpha\sigma]^{1/2})/2$. This is the value of w where applicants choose a cutoff $\tilde{a} = w$. We would expect that the firm would not choose $w > \bar{w}$ because it would then run the risk of earning negative profits *ex post*. As this intuition suggests, the derivative of the profit function is negative when evaluated at $\bar{w}$. Specifically, the second term simplifies to $(3-2z)(\alpha-\tilde{a})(\partial z/\partial w)$, and it is then easy to show that the first and second terms cancel, implying $\partial(3\Pi^{os})/\partial w = -3z(2-z) < 0$.

We now show the following sufficient conditions for the second derivative of the profit function to be negative: (i) $w < \bar{w}$ and (ii) $w > \underline{w} \approx 0.65\sigma$. Differentiating the profit function a second time, we obtain

$$\begin{aligned}\frac{\partial^2(3\Pi^{os})}{\partial w^2} =& -6(1-z)\frac{\partial z}{\partial w} \\ &+ [\alpha(3-2z)+\tilde{a}(3-4z)-6w(1-z)]\frac{\partial^2 z}{\partial w^2} \\ &- [2\alpha+4\tilde{a}-6w]\left(\frac{\partial z}{\partial w}\right)^2 \\ &+ [z(3-2z)]\frac{\partial^2 \tilde{a}}{\partial w^2} \\ &+ 2(3-4z)\frac{\partial z}{\partial w}\frac{\partial \tilde{a}}{\partial w}.\end{aligned}$$

The first term is clearly negative. Consider the second term. Assume $w = \tilde{a}$. Then this term becomes $(3-2z)(\alpha-\tilde{a})(\partial^2 z/\partial w^2) < 0$; it is negative *a fortiori* for $w < \tilde{a}$. Condition (i) similarly guarantees that the third term is negative. Now, using the facts that $\partial\tilde{a}/\partial w = -((w+\sigma)/2)(\partial^2\tilde{a}/\partial w^2)$ and $\sigma/(w+\sigma) = 1-z$, combine the last two terms to obtain

$$[z(3-2z)-(3-4z)(1-z)]\,\frac{\partial^2\tilde{a}}{\partial w^2} = -(6z^2-10z+3)\frac{\partial^2\tilde{a}}{\partial w^2}$$

which is negative if $6z^2 - 10z + 3 > 0$. The larger root exceeds one; the smaller root equals approximately 0.39. In order to ensure $z > 0.39$, we require $w > 0.65\sigma$. We have thus shown that the (continuous) profit function is increasing at $w = 0$, concave on the range $[\underline{w}, \bar{w}]$, and decreasing at $w = \bar{w}$ and thereafter; it thus has a maximum in $[0, \bar{w}]$. (While we have not characterised the function below $\underline{w}$, simulations suggest that it is in fact globally concave.)

Proof of Lemma 1

Both results follow immediately from Property 3 of the $\pi_H()$ and $\pi_L()$ functions. ■

Proof of Proposition 1

The proof proceeds as follows. For each region in turn, we first confirm that the proposed strategies are indeed a Nash equilibrium, and then show uniqueness by ruling out other strategies.

1. In the first region, suppose that any candidate, when considering whether to apply, conjectures that $\bar{z}_H = 2w/(q(w+\sigma)) < 1$ and $\bar{z}_L = 0$. We must confirm that these strategies represent a best response for a high-ability and a low-ability applicant respectively.

(a) Consider a high-ability applicant. She will get the job with probability $(1-q) + q[\bar{z}_H/2 + (1-\bar{z}_H)] = (2-q\bar{z}_H)/2$. For her to optimally choose a mixed strategy, it must be the case that

$$\pi_H(\bar{z}_H,0)w-(1-\pi_H(\bar{z}_H,0))\sigma=0$$
$$\Rightarrow \frac{w(2-q\bar{z}_H)}{2}=\left(1-\frac{(2-q\bar{z}_H)}{2}\right)\sigma$$
$$\Rightarrow \bar{z}_H=\frac{2w}{q(w+\sigma)}=z_H^*.$$

(b) Now consider a low-ability candidate. It follows immediately from Lemma 1 that she will not apply.

(c) To prove uniqueness in this region, observe that in this region of the parameter space, $\pi_H(1,0)w-(1-\pi_H(1,0))\sigma<0$. Thus, when low-ability agents do not apply, it is not an equilibrium for high-ability agents to apply with probability one. From Property 2 above, if low-ability agents were to apply with positive probability, it would *a fortiori* not be an equilibrium strategy for high-ability agents to apply with probability one. Thus it is never an equilibrium for high-ability agents to apply with probability one, and so it follows from Lemma 1 that it is never an equilibrium for low-ability agents to apply with positive probability.

2. In the second region, suppose that any agent, when considering whether to apply, conjectures that $\bar{z}_H=1$ and $\bar{z}_L=0$.

(a) Consider a high-ability applicant. She will get the job with probability $1-q+q/2=(2-q)/2$. Thus her expected payoff is

$$\frac{w(2-q)}{2}-\sigma\frac{q}{2}\geq 0$$

If $w<\sigma q/(2-q)$, then this implies $z_H^*=1$. If $w=\sigma q/(2-q)$ the agent is indifferent among all $z_H^*\in[0,1]$, and $z_H^*=1$ is still a best response.

(b) Consider a low-ability applicant. She will get the job with probability $1-q$. (Recall the conjecture is $\bar{z}_L=0$.) Thus her expected payoff is

$$w(1-q)-\sigma q\leq 0$$

If $w<\sigma q/(1-q)$, then this implies $z_L^*=0$. If $w=\sigma q/(1-q)$ the agent is indifferent among all $z_L^*\in[0,1]$, and $z_L^*=0$ is still a best response.

(c) To prove uniqueness in this region, note that the second inequality in the definition of the parameter space implies that $\pi_L(1,0)w-(1-\pi_L(1,0))\sigma\leq 0$. Consider first the case of a strict inequality: $\pi_L(1,0)w-(1-\pi_L(1,0))\sigma<0$. Then if high-ability agents are applying with probability one, it can never be an equilibrium strategy for a low-ability agent to apply with positive probability. But we already know that if high-ability agents apply with probability less than one, it also can never be an equilibrium strategy for a low-ability agent to apply with positive probability. If the condition holds with equality, then a low-ability agent is indifferent among all $z_L^* \in [0, 1]$, given the conjecture that $\bar{z}_L = 0$. But for any $\bar{z}_L > 0$, the unique best response of a low-ability agent is $z_L^* = 0$.

Similarly, the first inequality in the definition of the parameter place implies $\pi_H(1, 0)w - (1 - \pi_H(1, 0))\sigma \geq 0$. If this inequality is strict, it is not an equilibrium for low-ability agents not to apply and for high-ability agents to apply with probability less than one. If the expression holds with equality, then a high-ability agent is indifferent among all $z_H^* \in [0, 1]$, given the conjecture that $\bar{z}_H = 1$. But for any $\bar{z}_H < 1$, the unique best response of a high-ability agent is $z_H^* = 1$.

3. In the third region, suppose that any agent, when considering whether to apply, conjectures that $\bar{z}_H = 1$ and $0 < \bar{z}_L = 2\left[1 - \sigma/((w + \sigma)(1 - q))\right] < 1$.

(a) Consider a low-ability applicant. She will get the job with probability $(1-q)(1-\bar{z}_L+\bar{z}_L/2) = (1 - q)(1 - \bar{z}_L/2)$. For her to optimally choose a mixed strategy, it must be the case that

$$w(1-q)\frac{2-\bar{z}_L}{2} - \left[1 - \frac{(1-q)(2-\bar{z}_L)}{2}\right]\sigma = 0$$
$$\Rightarrow \bar{z}_L = 2\left[1 - \frac{\sigma}{(w+\sigma)(1-q)}\right]$$
$$= z_L^* > 0.$$

(b) Now consider a high-ability candidate. It follows immediately from Lemma 1 that she will apply with probability 1.

(c) To prove uniqueness in this region, note that the first inequality in the definition of the parameter space implies that $\pi_L(1, 0)w - (1 - \pi_L(1, 0))\sigma > 0$. This means that it can never be an equilibrium strategy for a low-ability agent to apply with probability zero (and thus, from Lemma 1, also that it can never be an equilibrium strategy for a high-ability agent to apply with probability less than one). The second inequality in the definition of the parameter space implies $\pi_L(1, 1)w - (1 - \pi_L(1, 1))\sigma < 0$. This means that it is not an equilibrium strategy for low-ability agents to apply with probability one.

4. In the fourth region, suppose that any candidate, when considering whether to apply, conjectures that $\bar{z}_H = 1$ and $\bar{z}_L = 1$.

(a) Consider a low-ability applicant. She will get the job with probability $(1 - q)/2$. Thus her expected payoff is

$$w\left[\frac{1-q}{2}\right]-\sigma\left[\frac{1+q}{2}\right]\geq 0$$

If $\sigma(1+q)/(1-q)<w$, then this implies $z_L^*=1$. If $\sigma(1+q)/(1-q)=w$ the agent is indifferent among all $z_L^*\in[0,1]$, and $z_L^*=1$ is still a best response.

(b) Consider a high-ability applicant. She will get the job with probability $1-q+q/2=(2-q)/2$. Thus her expected payoff is

$$\frac{w(2-q)}{2}-\sigma\frac{q}{2}>0$$

$$\Rightarrow z_H=1.$$

(c) To prove uniqueness in this region, note that the inequality in the definition implies $\pi_L(1,1)w-(1-\pi_L(1,1))\sigma\geq 0$. Consider first the case of a strict inequality: $\pi_L(1,1)w-(1-\pi_L(1,1))\sigma>0$. Then low-ability candidates always apply with probability one, and from Lemma 1, high-ability candidates also always apply with probability one. If the condition holds with equality, then a low-ability agent is indifferent among all $z_L^*\in[0,1]$, given the conjecture that $\bar{z}_L=1$. But for any $\bar{z}_L<1$, the unique best response of a low-ability agent is $z_L^*=1$. ■

Proof of Lemma 2

If the firm chooses a separating strategy, it will choose $w_{os}\in(0,\sigma q/(2-q)]$. This follows from Proposition 1, together with the observation that $w_{os}=\sigma q/(2-q)\Rightarrow z_H=1$; setting a higher salary would increase the firm's costs while leaving expected revenue unchanged, and therefore could not be an optimal strategy.

Differentiating the profit function yields

$$\frac{\partial\Pi^{os}}{\partial w_{os}}=\frac{4\sigma}{(w_{os}+\sigma)^3}\left[\sigma\alpha-w(2\sigma+\alpha)\right].$$

The second derivative of the profit function is negative, and thus the profit function has a unique maximum.

Setting the derivative equal to zero we see that, at an interior solution, $w_{os}=\sigma\alpha/(2\sigma+\alpha)$. If $\alpha<\sigma q/(1-q)$, then it is straightforward to show that this interior solution satisfies $w_{os}<\sigma q/(2-q)$, which implies $z_H^*<1$. If $\alpha\geq\sigma q/(1-q)$, then the profit function attains its maximum at the corner solution $w_{os}=\sigma q/(2-q)$, where $z_H^*=1$. It is easily verified that in either case $\alpha>w_{os}$, so the firm does indeed wish to hire at the posted salary. ■

Proof of Lemma 3

Suppose not. That is assume that the firm, in equilibrium, chooses the salary that maximises Π^{os}, and that this solution is an interior maximum. Then we know that $\alpha \leq \sigma q/(2-q) \Rightarrow q \geq \alpha/(\sigma+\alpha)$ and $\Pi^{os} = \alpha^2/(\sigma+\alpha)$.

Now, recall that the firm's profit function under secret search is: $\Pi^{ss} = q\alpha + (1-q)$, which is increasing in q.

$$\begin{aligned}
\Pi^{ss} &= q\alpha + (1-q) \\
&\geq \frac{\alpha}{\sigma+\alpha}\alpha + (1 - \frac{\alpha}{\sigma+\alpha}) \\
&= \frac{\alpha^2+\sigma}{\sigma+\alpha} \\
&> \Pi^{os}.
\end{aligned}$$

Therefore this cannot be an equilibrium strategy for the firm.

Proof of the Comparative Static Results

When the firm adopts a pooling strategy, profits are given by

$$\Pi^{op} = \bar{\psi}[\alpha - w] + (1-\bar{\psi})\mu(w)[1-w].$$

Let

$$f(w, q; \sigma) = \frac{\partial \Pi^{op}}{\partial w} = (1-\bar{\psi})(1-w)\mu'(w) - [\bar{\psi} + (1-\bar{\psi})\mu(w_{op})] = 0$$

so the optimal salary at a pooling equilibrium, w^*_{op}, is implicitly defined by

$$f(w^*_{op}, q; \sigma) = (1-w^*_{op})\mu'(w^*_{op}) - \mu(w^*_{op}) - \frac{\bar{\psi}}{1-\bar{\psi}} = 0.$$

Now, for any salary consistent with a pooling equilibrium ($w \in (\sigma q/(1-q), \sigma(1+q)/(1-q)]$),

$$\begin{aligned}
\mu(w) &= \frac{4\sigma[w - q(w+\sigma)]}{(w+\sigma)^2(1-q)^2} > 0 \\
\mu'(w) &= \frac{4\sigma[(\sigma - w) + q(w+\sigma)]}{(w+\sigma)^3(1-q)^2} \geq 0 \\
\mu''(w) &= \frac{4\sigma[(3\sigma - w) + q(w+\sigma)]}{(w+\sigma)^4(1-q)^2} > 0
\end{aligned}$$

Since

$$\frac{\partial f(w^*_{op}, q; \sigma)}{\partial w^*_{op}} = -\mu'(w^*_{op}) - w^*_{op}\mu''(w^*_{op}) < 0,$$

the profit function is strictly concave with a unique maximum at $w^*_{op} \in (\sigma q/(1-q), \sigma(1+q)/(1-q))$. (Because $f(\sigma(1+q)/(1-q), q; \sigma) < 0$, it is never optimal for the firm to set the salary at the level where all agents apply with probability one.) Note that $f(1, q; \sigma) < 0$; this implies that the firm sets $w^*_{op} < 1$, ensuring that the firm earns positive profits even if it hires a low-ability applicant.

We now consider each of the results, in turn.

w_{op} is independent of α.

This follows immediately from the fact that the first-order condition that defines w_{op} does not contain α.

w_{op} is increasing in σ.

To show that w is increasing in σ, we need to show that $\partial f(w^*_{op}, q, \sigma)/\partial\sigma > 0$. Here, it is easier to begin from the definition of $\mu() = z_L(2 - z_L)$. Now

$$\begin{aligned}
&\frac{\partial f(w^*_{op}, q, \sigma)}{\partial \sigma} = \left[(1 - w^*_{op})\frac{\partial \mu'(w^*_{op})}{\partial z_L} - \frac{\partial \mu(w^*_{op})}{\partial \partial z_L}\right]\left(\frac{\partial z_L}{\partial \sigma}\right) \\
&= \left[\frac{2(1 - w^*_{op})}{(w_{op} + \sigma)^2}\left[(w_{op} + \sigma)(2z_L - 3) - (1 - z_L)(2 - z_L)\right] - 2(1 - z_L)\right]\left(\frac{\partial z_L}{\partial \sigma}\right) \\
&> 0
\end{aligned}$$

where we have used the fact that $\mu'() = 2(1 - z_L)(2 - z_L)/(w + \sigma)$.

w_{op} is increasing in q.

By the implicit function theorem

$$\frac{\partial w^*_{op}}{\partial q} = -\frac{\frac{\partial f(w^*_{op}, q, \sigma)}{\partial q}}{\frac{\partial f(w^*_{op}, q, \sigma)}{\partial w^*_{op}}}$$

so we need to show that $\partial f(w^*_{op}, q, \sigma)/\partial q > 0$. Differentiation yields

$$\begin{aligned}
&\frac{\partial f(w^*_{op}, q; \sigma)}{\partial q} = \left[\frac{1}{1 - q}\right]^3 \\
&\times \left[\left(\frac{4\sigma}{(w^*_{op} + \sigma)^3}\right)\left[\sigma^2 - w^*_{op} + 3\sigma(1 - w^*_{op}) + q(w^*_{op} + \sigma)(1 + \sigma)\right] - 2\right]
\end{aligned}$$

This expression is increasing in q and decreasing in w. We first show that it is positive when evaluated at $q = 0$. Evaluating the first-order condition at $q = 0$ yields $w^*_{op} = \sigma/(1 + 2\sigma)$. Using this, we obtain

$$\begin{aligned}\frac{\partial f(w^*_{op}, 0, \sigma)}{\partial q} &= \frac{4\sigma}{(w^*_{op}+\sigma)^3}\left[\sigma^2 - w^*_{op} + 3\sigma(1-w^*_{op})\right] - 2 \\ &= \frac{(1+2\sigma)^2}{\sigma(1+\sigma)} - 2 \\ &> 0.\end{aligned}$$

Now we prove that $\partial f(w^*_{op}, q, \sigma)/\partial q$ cannot ever equal zero. The proof proceeds by contradiction. Assume that $\exists \hat{q} \in (0,1)$ such that $\partial f(w^*_{op}, \hat{q}, \sigma)/\partial q = 0$ and $\partial f(w^*_{op}, q, \sigma)/\partial q > 0 \forall q < \hat{q}$. By continuity, $\partial^2 f(w^*_{op}, \hat{q}; \sigma)/\partial q^2 \leq 0$. But $\partial f(w^*_{op}, \hat{q}; \sigma)/\partial q = 0 \Rightarrow \partial w^*_{op}/\partial q|_{q=\hat{q}} = 0 \Rightarrow d^2 f(w^*_{op}, \hat{q}; \sigma)/dq^2 = \partial^2 f(w^*_{op}, \hat{q}; \sigma)/\partial q^2 > 0$, so we have a contradiction. Hence the salary in the pooling equilibrium is increasing in q.

w_{os} is independent of α.

This and the following two results follow immediately from the fact that, in the separating regime, $w_{os} = \sigma q/(2-q)$ in equilibrium.

w_{os} is increasing in σ.

See above.

w_{os} is increasing in q.

See above.

z_H is independent of α.

This and the following two results follow immediately from the fact that, in both the separating and the pooling regime, the firm chooses the salary such that $z_H = 1$ in equilibrium.

z_H is independent of σ.

See above.

z_H is independent of q.

See above.

z_L is independent of α.

In the separating regime, $z_L = 0$ by definition. In the pooling regime, $z_L = 2[1 - \sigma/((w+\sigma)(1-q))]$. Given that the salary in the pooling regime is independent of α, the result follows.

z_L is independent of σ in the separating regime and decreasing in σ in the pooling regime.

The result is immediate for the separating regime. We do not have an analytical proof for the pooling regime, but the result is supported by extensive simulations.

z_L is independent of q in the separating regime and decreasing in q in the pooling regime.

The result is immediate for the separating regime. We do not have an analytical proof for the pooling regime, but the result is supported by extensive simulations.

Π^{os} *is increasing in* α.

Starting from equilibrium, an increase in α directly increases the firm's profits. Thus we know that the firm could hold the salary fixed and enjoy higher profits. It follows that profits increase *a fortiori* when the firm optimally adjusts the salary.

Π^{os} *is decreasing in* σ.

In the separating regime, $\Pi^{os} = \bar{\psi}\alpha - \sigma q^2$, where $\bar{\psi} = q(2-q)$. The result is immediate.

Π^{os} *is strictly increasing in* q.

Differentiating the profit function, we obtain $\partial\Pi/\partial q = 2(1-q)[\alpha - \sigma q/(1-q)] > 0$ (given that we are, by assumption, in the separating regime).

Π^{op} *is increasing in* α.

Starting from equilibrium, an increase in α directly increases the firm's profits. Thus we know that the firm could hold the salary fixed and enjoy higher profits. It follows that profits increase *a fortiori* when the firm optimally adjusts the salary.

Π^{op} *is decreasing in* σ.

Expected profit in the pooling equilibrium is strictly decreasing in σ. This follows immediately from the fact that profit is strictly increasing in $\mu()$, $\mu()$ is strictly increasing in z_L, and z_L is strictly decreasing in σ. The effect of changes in the salary can be ignored because of the envelope theorem.

Π^{op} *is increasing or decreasing in* q.

The effects of increases in q on expected profit in the pooling equilibrium are ambiguous. Specifically

$$\frac{\partial \Pi^{op}}{\partial q} = [(\alpha - w^*_{op}) - \mu(w^*_{op})(1-w^*_{op})]\frac{\partial\bar{\psi}}{\partial q} + (1-\bar{\psi})(1-w^*_{op})\frac{\partial\mu(w*)}{\partial q}$$

(By the envelope theorem, we can again ignore the effects of change in q on the salary.) The first term is positive, reflecting the fact that increases in q shift the probability mass to hiring a high-ability applicant rather than a low-ability applicant. As high-ability applicants are more productive, this increases expected profits. But the second term captures the fact that increases in q make it less likely that low-ability candidates will apply ($\partial\mu(w^*_{op})/\partial q < 0$). The first term depends on the value of α while the second term does not. After some algebra, we can write this expression as

$$\frac{\partial \Pi^{op}}{\partial q} = 2(1-q)\left[(\alpha-1) - (1-z_L)(1-w^*_{op})\right]$$

which can be greater or less than 0 depending on the value of α. ■

Proof of Proposition 2

Under a separating strategy, $z_L = 0 \Rightarrow \frac{\partial z_L}{\partial w_{pp}} = 0$. Thus

$$\frac{\partial W}{\partial w_{pp}} = \left[2q\alpha(1 - qz_H) - 2q^2\sigma z_H\right] \frac{\partial z_H}{\partial w_{pp}} = 0$$

If $\alpha < \sigma q/(1-q)$, then this equation has an interior solution:

$$\Rightarrow z_H = \frac{\alpha}{q(\alpha + \sigma)}$$

Given $z_H = 2w/(q(w+\sigma))$, it follows that

$$w_{pp} = \frac{\alpha\sigma}{\alpha + 2\sigma}$$

which is the same solution as for the firm's problem.

If $\alpha \geq \sigma q/(1-q)$ then the first-order condition in part 1 is satisfied at the corner solution where $\partial z_H/\partial w_{pp} = 0$. This requires setting $w_{pp} \geq \sigma q/(1-q)$. In the case where $w_{pp} = \sigma q/(1-q)$, the solution is identical to that from the firm's problem. The planner is indifferent between this salary and any higher salary consistent with a separating strategy: increases in the salary represent a redistribution from the firm to the employee, but do not affect behaviour or expected output.

Under a pooling strategy, $z_H = 1, \partial z_H/\partial w_{pp} = 0$, $\psi = \bar{\psi} = q(2-q)$, and

$$\frac{\partial W}{\partial w_{pp}} = \left[2(1-\bar{\psi})(1-z_L) - 2(1-q)\sigma[q + (1-q)z_L)]\right] \frac{\partial z_L}{\partial w_{pp}} = 0$$

At an interior solution, this implies

$$z_L = \frac{1 - q(1+\sigma)}{(1-q)(1+\sigma)}$$

Using the fact that $z_L = 2[1 - \sigma/((w+\sigma)(1-q))]$ we obtain (after some algebra)

$$w^*_{pp} = \frac{\sigma(1 + q + q\sigma)}{1 - q - q\sigma + 2\sigma}$$

For the firm's problem, we do not have a closed form solution for the salary, so we cannot make a direct comparison between the planner's solution and the equilibrium solution. Recall, however, that

$$
\begin{aligned}
\frac{\partial \Pi^{op}}{\partial w} =& (1-\bar{\psi})(1-w)\mu'(w) - [\bar{\psi} + (1-\bar{\psi})\mu(w)] \\
=& (1-\bar{\psi})[(1-w)\mu'(w) - \mu(w)] - \bar{\psi}
\end{aligned}
$$

It can be shown that

$$
\begin{aligned}
1 - w^*_{pp} &= \frac{(1+\sigma)[(1-q)(1+\sigma)-\sigma]}{(1-q)(1+\sigma)+\sigma} \\
\mu(w^*_{pp}) &= \frac{[(1-q)(1+\sigma)]^2 - \sigma^2}{[(1-q)(1+\sigma)]^2} \\
\mu'(w^*_{pp}) &= \frac{[(1-q)(1+\sigma)+\sigma)]^2}{(1+\sigma)[(1-q)(1+\sigma)]^2}
\end{aligned}
$$

from which we can show that, evaluated at w^*_{pp},

$$
\frac{\partial \Pi^{op}}{\partial w} = -\bar{\psi} < 0.
$$

Thus, starting from the planner's solution for the salary, a reduction in the salary increases the firm's expected profits. From the concavity and continuity of the profit function it follows that $w^*_{op} < w^*_{pp}$.

The logic is straightforward. When the planner is choosing the salary, she trades off the marginal benefit of inducing more low-ability agents to apply against the marginal cost in terms of expected disutility of rejection of those agents. At the optimum, the planner equalises these marginal costs, ignoring the fact that the salary is also paid to high-ability agents. The planner does not care about the fact that the salary must also be paid to high-ability agents, because the salary has no effect at the margin on the behaviour of those agents. The firm, however, treats the salary paid to high-ability agents as an additional cost.

In Figure 3 (which shows different strategy choices in q-α space), consider any point on the boundary between pooling and secret search. By definition, the firm is indifferent between the two strategies ({secret search with $w = w^*_{ss} = 0$, open pooling with $w = w^*_{op}$}). At this point, it follows from part (3) above that there exists a pooling strategy with $w_{pp} > w^*_{op}$ which yields strictly higher welfare. By continuity, there exist neighbouring points where the firm would choose secret search but where pooling yields higher welfare. Exactly the same argument holds on the boundary between pooling and separating strategies. ■

References

[1] Albrecht, J., P. Gautier and S. Vroman, "Equilibrium Directed Search with Multiple Applications," *Review of Economic Studies*, 73 (2006), 869-891.

[2] Alonso, R., "Recruitment and Selection in Organizations," *USC Marshall School of Business Research Paper Series*, University of Southern California (2014).

[3] Auster, S. and P. Gottardi, "Competing Mechanisms in Markets for Lemons," *IGIER Working Paper* 568 (2017).

[4] Basov, S., I. King and L. Uren, "Worker Heterogeneity, the Job Finding Rate, and Technical Change," *European Economic Review,* 70 (2014), 159-177.

[5] Bebchuk, L. and J. Fried, "Executive Compensation as an Agency Problem," *Journal of Economic Perspectives*, 17, 3 (2003), 71-92.

[6] Cai, X., P. Gautier, and R. Wolthoff, "Search Frictions, Competing Mechanisms and Optimal Market Segmentation," *Journal of Economic Theory*, 169 (2017), 453-473.

[7] Edmans, A. and X. Gabaix, "Executive Compensation: A Modern Primer," *Journal of Economic Literature*, 54(4) (2016), 1232-1287.

[8] Eeckhout, J. and P. Kircher, "Sorting versus Screening: Search Frictions and Competing Mechanisms," *Journal of Economic Theory*, 145 (2010), 1354-1385.

[9] Galenianos, M. and P. Kircher, "Directed Search with Multiple Job Applications," *Journal of Economic Theory*, 144 (2009), 445-471.

[10] Guerrieri, V., R. Shimer, and R. Wright, "Adverse Selection in Competitive Search Equilibrium," *Econometrica*, 78 (2010), 1823-1862.

[11] Julien, B., J. Kennes and I. King, "Bidding for Labor," *Review of Economic Dynamics*, 3 (2000), 619-649.

[12] Kaplan, S., M. Klebanov, and M. Sorenson, "Which CEO Characteristics and Abilities Matter?" *The Journal of Finance*, 47,3 (2012), 973-1007.

[13] Khurana, R., "Leadership and the Social Construction of Charisma," in J. P. Doh and S. A. Stumpf, eds., *Handbook on Responsible Leadership and Governance in Global Business*, (Cheltenham: Edward Elgar, 2005), 112-136.

[14] McAfee, P. "Mechanism Design by Competing Sellers," *Econometrica*, 61 (1993), 1281-1312.

[15] Oyer, P. and S. Schaefer, "Personnel Economics: Hiring and Incentives," in O. Ashenfelter and D. Card, *Handbook of Labor Economics*, Elsevier (2010).

[16] Peters, M., and S. Severinov, "Competition Among Sellers Who Offer Auctions Instead of Prices," *Journal of Economic Theory*, 75 (1997), 141-179.

[17] Rajgopal, S., D. Taylor and M. Venkatachalam, "Frictions in the CEO Labor Market: The Role of Talent Agents in CEO Compensation," *Contemporary Accounting Research*, 29,1 (2012), 119-151.

[18] Rees, A., "Information Networks in Labor Markets," *American Economic Review*, 56 (1966), 559-566.

[19] Rogerson, R., R. Shimer and R. Wright, "Search Theoretic Models of the Labor Market: A Survey," *Journal of Economic Literature*, XLIII (2005), 959-9888.

[20] Salop, J. and S. Salop, "Self-selection and Turnover in the Labor Market," *Quarterly Journal of Economics*, 90 (1976), 619-627.

[21] Theune, M. and H. Tiede, "Confidential Searches for Chief Academic Officers," *Academe*100.6 (2014), 31-33.

[22] van Ours, J. and G. Ridder, "Vacancies and the Recruitment of New Employees," *Journal of Labor Economics*, 10 (1992), 138-155.

[23] Vinkenburg, C., P. Jansen, N. Dries and R. Pepermans, "Arena: A Critical Conceptual Framework of Top Management Selection," *Group & Organization Management*, 39(1) (2014), 33-68.

[24] Wolthoff, R. "Applications and Interviews: A Structural Analysis of Two-Sided Simultaneous Search," University of Toronto manuscript (2012).

[25] Wright, R., P. Kircher, B. Julien and V. Guerrieri, "Directed Search: A Guided Tour," *NBER Working Paper* 23884 (2017).

[26] Zajac, E., "CEO Selection, Succession, Compensation and Firm Performance: A Theoretical Integration and Empirical Analysis," *Strategic Management Journal*, 11, 3 (1990), 217-230.