Reputation, Competition, and Lies in Labor Market Recommendations *

ODILON CÂMARA† NAN JIA‡ JOSEPH RAFFIEE§

Marshall School of Business

University of Southern California

September 26, 2022

*Authors are listed alphabetically. We thank the editor, associate editor, and referees for their valuable suggestions.

†Department of Finance and Business Economics, 701 Exposition Blvd., HOH 822, Los Angeles, CA, 90089, ocamara@marshall.usc.edu

‡Department of Management and Organization, 701 Exposition Blvd., HOH 518, Los Angeles, CA, 90089, nan.jia@marshall.usc.edu

§Department of Management and Organization, 701 Exposition Blvd., HOH 512, Los Angeles, CA, 90089, joe.raffiee@marshall.usc.edu

Abstract

We examine strategic communication in labor market recommendations. Our formal model features two-sided asymmetric information: an adviser has private information about his own preference bias for a focal candidate and a signal of the quality of this candidate, while the hiring firm has private information about the quality of an alternative candidate. The adviser can choose whether to recommend his focal candidate to the firm. If he recommends and the firm hires the candidate, then the adviser pays a reputational cost (receives a reputation boost) if the firm later learns that the hire has low quality (high quality). Our main results describe how the equilibrium behavior of advisers (lying choices) and firms (hiring choices) depend on the intricate interplay between preference biases, reputation, lying costs, and the hiring firm's labor market strength (access to alternative candidates with higher quality). We show that the equilibrium features assortative matching: advisers with a higher (lower) reputation choose to lie less (more) and, consequently, their candidates are more likely to be hired by firms with strong (weak) access to high-skilled outside candidates. Two equilibrium forces create a "rich get richer" effect. First, advisers choose to lie less to hiring firms with access to better top-candidates, further benefiting those firms. Second, advisers with a higher (lower) reputation choose to lie less (more), which increases (decreases) their future reputation, creating a "reputation trap." We discuss the implications of our model for hiring strategy, referral systems, and the ability to accrue and sustain human capital-based competitive advantages.

1 Introduction

Strategy scholars and personnel economists have long acknowledged the importance of making quality hiring decisions (Bidwell 2011; Groysberg, Lee and Nanda 2008; Singh and Agrawal 2011; Starr, Ganco and Campbell 2018). Hiring is a key pathway through which firms build and advance their human capital profiles, which, in turn, affect firm performance (e.g., Gambardella, Panico and Valentini 2015; Haltiwanger, Lane and Spletzer 1999; Hatch and Dyer 2004). This intuition is well-captured by Oyer and Schaefer (2010: 3): "hiring the right employee is potentially as important or more so than motivating the employee to take the right action after the employee has been hired". However, making correct hiring decisions is not straightforward. Central challenges that hinder firms from making the right hiring decisions include information asymmetries which make it tough for firms to accurately assess applicant quality prior to hire and search costs which prevent firms from observing all potential applicants (Jovanovic 1979; Spence 1973; Stigler 1961, 1962). To overcome these challenges, a common strategy firms employ is to solicitate recommendations from their workforce (Brown, Setren and Topa 2016; Burks et al. 2015; Castilla 2005) and/or other sources (Castilla and Rissing 2018; Abel, Burger, Piraino, 2020).[1]

For referrals to lead to better hires/matches (Fernandez and Weinberg 1997; Montgomery 1991; Topa 2011; Dariel et al. 2021), two key assumptions are essential: (1) that recommenders are more informed about candidates than the hiring firm (Greenwald 1986), and (2) that recommenders send truthful recommendations in good faith (i.e., recommend high-ability candidates and do not recommend low-ability candidates).[2] However, strategic recommenders may be tempted to send misleading or false recommendations.[3] Recognizing this challenge, the literature has explored how firms can use incentives to elicit truthful recommendations. One direct mechanism would be for the firm to offer a formal payment for recommendations (Burks et al., 2015). The literature has also studied indirect mechanisms

[1]Topa (2011) estimates that approximately half of all jobs in the United States are filled via employee referrals. Similarly, Castilla and Rissing (2018) suggest that informal recommendations and endorsements made by outside parties, not just employees, are likely to be quite prevalent, although the nature of the phenomena makes the data trail elusive.

[2]Utilizing referrals in the hiring process may also be beneficial if referred workers put forth more effort (e.g., Heath, 2018) and/or work more productively when paired with the referring worker (e.g., Pallais and Sands, 2016).

[3]For example, Saloner (1985), Carmichael (1988), Friebel and Raith (2004), and Ekinci (2016) theoretically explore how misaligned incentives may lead to opportunistic behavior in the context of hiring. However, as highlighted by several scholars, this issue has been largely overlooked in the existing theoretical work, which largely assumes that the incentives of the employee "are perfectly aligned with those of the firm in hiring decisions" (Ekinci, 2016: 690), as well as empirical work where researchers often "assume that the provider and firm's incentives are aligned without having the data to validate the assumption" (Heath, 2018: 1695). Thus, the literature generally assumes that "referees truthfully report the information at their disposal" (Farchamps and Moradi, 2015: 716).

as incentives, such as competition (Saloner 1985) and signaling (Ekinci 2016).

In this paper, we focus on an alternative mechanism in which the recommender and the firm engage in a form of relational contract to provide a framework to think about key factors affecting the truthfulness of referrals and hiring firms' decisions.[4] In our formal model, a recommender or "adviser" has a relationship with a hiring firm or "decision maker." The adviser privately observes if his focal candidate is of high or low quality, and then chooses whether or not to recommend his candidate to the hiring firm. We assume that the adviser captures a private benefit if the focal candidate is hired, thereby introducing bias towards inflating the recommendation. The firm faces a choice of hiring the adviser's candidate or hiring an alternative candidate from the labor market. The key endogenous variable ruling this relationship is the adviser's "reputation" which follows an intuitive dynamic rule: the adviser's reputation increases after a successful hire (recommended candidate is hired and proven to be of high quality) and decreases after an unsuccessful hire (recommended candidate is hired but found to be of low quality).[5]

Our main assumption is that the reputational cost of a lie increases with the adviser's current reputation. This dynamic can be interpreted as reputation is slow to build but quick to lose, something well-captured by Warren Buffet: "It takes 20 years to build a reputation and five minutes to ruin it. If you think about that, you'll do things differently." (Lowenstein, 1995: 111). An adviser with a low reputation has little to lose (i.e., a lower reputational cost) and hence will be more likely to lie. The opposite is true for an adviser with a high reputation. Hence, with a simple "slow to build, quick to lose" dynamic, reputation serves as a mechanism that summarizes past behavior and dictates future behavior. To simplify exposition, our benchmark model assumes that the adviser receives some exogenous payoff resulting from having a higher reputation. In Section 5 (Extensions), we discuss how this payoff can arise endogenously.

We start by considering the adviser's choice. When the candidate is of high quality, the adviser always chooses to recommend the candidate. If the high-quality candidate is hired, it will deliver the bias payoff and also increase the adviser's reputation. If the candidate is of low quality, however, the adviser faces a tradeoff: If the adviser recommends the low-

[4] See Malcomson (2013) for a review of the literature on relational contracts.

[5] Advisers in our model are not exogenously born as honest or dishonest, but rather their reputation is purely an endogenous construct akin to a mental account that keeps track of their past behavior. Reputation has bootstrap features: any adviser can increase his own reputation by sending truthful recommendations, which will generate more trust from the firm in the future. Reputation thus evolves organically through an adviser's own lying choices, which, in turn, affect the adviser's future lying behavior. Hence, our focus is on moral hazard (lying choices) and not on adverse selection. In the literature review, we discuss existing adverse selection models in which advisers are born with different types (for example, different skill levels in Ekinci 2016) and reputation is interpreted as the firm's belief about the adviser's exogenous type.

quality candidate and the candidate is hired, then the adviser receives the private benefit associated with his bias but at the cost of decreasing his reputation. In this case, the adviser's equilibrium strategy is defined by a simple threshold rule: he will lie and recommend a low-quality candidate if his personal bias is sufficiently high, and he will not recommend a low-quality candidate if his bias is low.

We next consider the firm's choice. The firm knows that the adviser has private information about the candidate's quality and that the adviser is likely biased; however, the firm does not know the extent of the bias. Thus, upon receiving a recommendation, the firm forms an opinion about the expected quality of the focal candidate, considering the possibility that the adviser may be acting untruthfully and recommending a low-quality candidate. Since the firm knows the adviser's current reputation, the firm can use it to correctly predict the probability that the adviser is lying. The firm will trust more (less) the recommendation of an adviser with a higher (lower) reputation and thus will be confident (skeptical) that the candidate is of high quality. Before making a hiring decision, the firm draws a candidate from the outside market, privately learns about his skill, and then makes a decision: the firm will hire the focal candidate if his expected quality is higher than the quality of the outside candidate, as is typical in hiring decisions where firms compare the qualities of candidates. Thus, the shape of the skill distribution of the firm's outside market plays a central role in equilibrium hiring decisions. We consider two types of firms – firms which have strong access to high-skilled outside candidates (which we call facing a "top-heavy" outside market) and firms who have weak access to high-skilled outside candidates (which we call facing a "bottom-heavy" outside market).[6]

While existing work has largely neglected the role of outside markets and alternative candidates, our parsimonious framework uncovers several insights relating the distribution of the outside market to equilibrium outcomes. First, our model generates an interesting assortative matching: firms facing a top-heavy outside market are more likely to hire a candidate from an adviser with a stronger reputation, whereas firms facing bottom-heavy outside market are more likely to hire a candidate from an adviser with a weaker reputation. The intuition is the following: when an adviser of weak reputation recommends his candidate, the firm highly discounts the recommendation and views him as an "average" candidate. A firm facing a top-heavy outside market is unlikely to hire this average candidate, while the opposite is true for a firm facing bottom-heavy firm. To have a better chance against competition in a top-heavy outside market, the adviser must be trustworthy (i.e., must have a strong reputation) so that his candidate with a good recommendation is indeed considered

[6]Our formal definition is that the skill distribution of a top-heavy (bottom-heavy) market features a strictly increasing (decreasing) probability density function.

high-skilled and can be hired.[7]

Second, the adviser will lie less to a firm when the firm gains better access to top-candidates. The intuition is that lying has a relatively smaller impact on the probability of hiring when a firm has better access to high-skilled outside candidates, as better access to high-skilled outside candidates means that competition for the job is stronger. Hence, given the relatively small benefit associated with lying, the adviser is less inclined to pay the lying cost and therefore lies less. Third, the adviser lies less when a firm gains a more informative signal about the outside candidate. With a more informative signal, the firm becomes better able to identify top outside candidates, which increases competition at the top and reduces the adviser's incentives to lie, as in our second insight.

Moreover, we define conditions such that advisers face a "reputation trap:" An adviser with a high reputation will be less likely to lie, resulting in an expected increase in his reputation; by contrast, an adviser with a low reputation will be more likely to lie, resulting in an expected decrease in his reputation. One implication of the reputation trap is that reputation differences across advisers are self-reinforcing, which can lead to sustained differences across time.

Together, our model demonstrates how the strategic choices of the adviser can amplify and sustain differences across firms. Conceptually, these insights are consistent with the notion that the "rich get richer" in markets for human capital (e.g., Bidwell et al. 2015). Prior research shows that higher-status firms are better able to attract high-quality candidates (Rider and Tan 2014), and our model demonstrates that those firms (which face top-heavy outside candidate market) are also given more reliable recommendations. By contrast, if a low-status firm starts to attract better bottom candidates, then it suffers from the drawback of receiving even less reliable recommendations. These factors may contribute to a virtuous circle enjoyed by advantaged firms (better candidates and more reliable information), which sustains differences across firms (Barney 1986).

1.1 Related literature and contributions

1.1.1 Literature on hiring based on labor market referrals

Our paper contributes to the growing literature on hiring through employee referrals (Fernandez, Castilla and Moore 2000) and endorsements (Castilla and Rissing 2018). While

[7] Put alternatively, if an adviser never lies, then he faces a "gamble": sometimes he provides a recommendation and thus the hiring probability goes up; sometimes he does not provide a recommendation and thus the hiring probability goes down. In contrast, when the adviser always lies, then there is no gamble: the firm fully discounts the recommendation and keeps its prior belief. For a firm with a top-heavy outside market, the gamble increases the overall hiring probability. For a firm with a bottom-heavy outside market, the gamble decreases the overall hiring probability.

there is consensus in the literature that advisers tend to have better information than firms about candidates (Pallais and Sands, 2016; Heath, 2018, Dustmann et al., 2016), incentives are often needed to motivate advisers to truthfully share their information (Beaman and Magruder, 2012). However, existing theoretical and empirical work in this area largely assumes that referrals are made honestly and in good faith (e.g., Greenwald 1986).

Recent papers have begun to explore how firms can offer a monetary (direct) incentive in exchange for recommendations, often in the form of a lump sum or flat payment contingent on continued employment of the referral candidate (Schlachter and Pieper, 2019). However, these payment systems do not necessarily solve the misalignment problem and potential for "referrer opportunism", including the possibility for recommenders to recommend low-quality workers (e.g., Fafchamps and Moradi, 2015). For example, Bond et al. (2018) finds that referral bonuses are associated with an increased likelihood that referrers will provide referrals for workers who are low-quality and Friebel et al. (2019) report a negative relationship between the size of a referral bonus and the quality of workers who are referred.

Indirect incentives could be used as alternative approaches to elicit truthful recommendations. In Saloner (1985), competition between two advisers results in advisers sending truthful recommendations, in the sense that the firm can make optimal hiring choices as if it had the same information as the advisers. In the career-concern model of Ekinci (2016), the main incentive is signaling within an adverse selection framework. In this model, advisers have an exogenous, unobservable skill level which is correlated to the skill of their candidates. Hence, an adviser can increase his own reputation by recommending a candidate with a higher skill, where reputation is defined as the firm's belief about the adviser's skill. A higher reputation implies a higher expected skill and, consequently, the firm will be willing to pay to the adviser a wage that is increasing in his reputation. There is no lying in the equilibrium of this model, in the sense that the firm can perfectly forecast which candidate is recommended by the adviser. In addition to no lies in equilibrium, these papers also differ from our model in that they do not study how the shape of the firm's outside market influences hiring decisions and how recommendations can vary over time with the adviser's reputation, two contributions of our study.

We contribute to this literature by studying a different incentive: we model the recommendation process as a form of reputational contract. Instead of reputation being about an exogenous characteristic of the adviser (hidden information or adverse selection[8]), reputation reflects the endogenous lying choices of this adviser over time (hidden action or moral

[8]One possible drawback of mechanism based on adverse selection and signaling is that advisers might lose their incentives to recommend high-skilled candidates once the firm learns more about the adviser's skill. If over time the firm perfectly learns the adviser's skill through his on-the-job performance, then the adviser has no incentive to recommend high-skilled candidates.

hazard).[9] Our equilibrium features lies, which we believe are an important empirical phenomenon, and sheds some light on the dynamics of reputation and lies over time, such as our reputational trap result. Moreover, we present several results connecting the shape of the outside market and the equilibrium lying/hiring behavior. By incorporating the role of the outside market into our model, we contribute to the broader literature on referrals by demonstrating that characteristics of the hiring firm, specifically the degree to which the firm has access to a top-heavy or bottom-heavy distribution of outside candidates, influences equilibrium outcomes. Existing work on referrals which has incorporated the outside market has been concerned with how referrals may impact aggregate market dynamics, market efficiency, and the formation of market wages (Dariel et al., 2021), or how the outside market wage puts limits on a firm's ability to punish workers who provide false recommendations (Heath, 2018). In contrast, we show how heterogeneity across firms with respect to firms' access to outside candidates – specifically, a top-heavy versus bottom-heavy distribution of outside candidates – leads to assortative matching where firms with access to a top-heavy outside market are more likely to hire candidates from high reputation advisers, whereas firms with access to a bottom-heavy outside market tend to hire candidates from low reputation advisers. Together these results strongly suggest a "rich get richer" dynamic in markets for human capital, as firms with access to better outside candidates receive more reliable information from advisers with high reputations. Hence, our insights are important for theoretical and empirical researchers, as well as practitioners, who should consider the economic forces that we highlight in this paper.

1.1.2 Literature on reputation and lying behavior

A comparison between our approach of formally modeling reputations and lying behavior with those in the general literature (including but not exclusive to the labor market referral context) is warranted. First, our consideration of the interaction between reputation cost and labor market competition is related to and extends the work of Hodler, Loertscher and Rohner (2014) who find that an intermediate lying cost maximizes the hiring probability. A key assumption in their model is that the decision maker has no private information about the alternative option. In contrast, we show that if the firm has private information about

[9] Bar-Isaac and Tadelis (2008) discuss how the term reputation has been used in the contexts of hidden information and hidden actions. They focus is on the reputation of firms. In the hidden information case, firms are exogenously endowed with a high-quality product or a low-quality product; reputation is then the buyer's belief that the firm is a high-quality firm. In the hidden action case, firms are the same ex-ante, but endogenously choose whether to produce the more costly high-quality product or the cheaper low-quality product; reputation is then an endogenous construct used to summarize past behavior and predict future behavior, providing incentives for the firm to produce the high-quality product. See also Mailath and Samuelson (2006, Chapter 15).

the outside candidate (which we believe is a realistic assumption in many scenarios), then this result would fundamentally change depending on the type of market competition. In a top-heavy market, a high lying cost (high reputation cost) maximizes the hiring probability, while the opposite is true for a bottom-heavy market. Hence, our results highlight the importance of considering the firm's private information and its market. Other models which have incorporated competition, such as Saloner (1985) who models competition between recommenders, are often less tractable to answer the kind of questions that we are raising in this paper.

Second, in developing our model, we also contribute to the broad literature on the economics of information, in particular signaling games (Spence 1973) and the study of strategic information transmission (strategic communication). Our main departure from the canonical signaling games[10] is that we allow the receiver to also have private information. Hence, we contribute to the branch of this literature which studies communication in two-sided asymmetric information games. Seidmann (1990) highlighted how private information by the receiver might change communication on cheap talk games; see Ishida and Shimizu (2019) for more references on the literature of cheap talk with a privately informed receiver. In most of these papers, receiver and sender have private information about the same issue (for example, both have private information about the same job market candidate). In contrast, in our model, the receiver has private information about the alternative candidate. Some papers (such as Seidmann 1990) do allow for the private information of sender and receiver to be uncorrelated; however, they consider more abstract games and have more limited applied results. In contrast, by focusing our model to study job market recommendation, we are able to uncover interesting results. For example, we show that the effects of reputation on hiring probabilities fundamentally depends on the quality distribution of outside candidates.

Third, a recommendation in our model is not a cheap talk message as in Crawford and Sobel (1982). Our model is closer to models of costly lies (Kartik 2009; Kartik, Ottaviani and Squintani 2007); see also Banks (1990) and Callander and Wilkie (2007) for models in which politicians pay a cost if they lie about their policy plans. Our lying cost takes the form of a reputation loss. Finally, our model is also related to Chung and Harbaugh (2019) who develop and test in the laboratory a recommendation game where the sender has a lying cost and a private bias. They find that subject behavior in a laboratory experiment is largely consistent with predictions from the cheap talk literature. Relevant to our setup, most subjects in their experiment exhibit some limited lying aversion and decision makers partially discounts a recommendation from biased experts.

The rest of the paper is organized as follows. Section 2 presents the model. Section 3

[10] We are following the definition in Kreps and Sobel (1994).

characterizes the equilibrium. The main results are presented in Section 4. In Section 5, we present extensions of the model. In Section 6, we discuss related applications, the empirical relevance of our results, and future research avenues. We present our conclusions in Section 7. All proofs are in the Appendix.

2 Model

Overview: Our model could be used to analyze many situations in which a *sender* provides information to persuade a *decision maker*. For concreteness, we interpret our model as an *adviser* (he) providing information to sway the hiring decision of a *firm* (it). The adviser has private information about the skill (low or high skill) of a particular candidate (called *focal candidate*) because they had a previous relation (e.g., the adviser was a mentor, teacher, supervisor or colleague). The firm has private information about the value of an alternative option (the expected skill of an *alternative candidate*). The adviser must choose whether to recommend his candidate to the firm or not. However, the adviser is biased: his payoff is higher if the firm hires his candidate. This creates an important conflict of interests because the adviser has an incentive to lie, that is, recommend his candidate even when his candidate has a low skill. Lying is costly, in the sense that recommending a low-skilled candidate will decrease the adviser's reputation if the firm hires the candidate and later observes his low skill on the job. Hence, an adviser who has a low-skilled candidate must decide whether he should lie (send a recommendation) and pay the cost to increase his candidate's chances of getting the job. To simplify the analysis, we present the model as a one period model and introduce reputation as a variable to capture relevant dynamic incentives.

Focal Candidate: Let $s \in \{s_L, s_H\}$ be the skill of the focal candidate, where $0 < s_L < s_H < 1$. With probability $p \in (0, 1)$ the focal candidate has a high skill s_H, and with probability $(1 - p)$ he has a low skill s_L. Therefore, the average skill of the focal candidate is $\hat{s} \equiv ps_H + (1 - p)s_L$. This probability distribution is common knowledge. The adviser privately knows the skill of his focal candidate, but the firm only knows its probability distribution.[11]

Recommendation: After observing the focal candidate's skill, the adviser must choose whether to recommend him to the firm or not. We denote this action as a message $m \in M \equiv \{NR, R\}$, where $m = R$ means sending a recommendation and $m = NR$ means not sending a recommendation. The firm wants the adviser to follow a simple *recommendation*

[11] We discuss extensions of the model in Section 5, including the case in which the adviser does not know the candidate's skill but privately observes a noisy signal of his skill.

standard: recommend the focal candidate if and only if he has a high skill. Therefore, we say that a recommendation is truthful if the adviser is recommending a high-skilled candidate, and a recommendation is untruthful if he is recommending a low-skilled candidate. The firm can only confirm the candidate's skill after it hires the candidate and observes his on-the-job performance. We call it a *successful hire* when the adviser recommends a high-skilled candidate and he is hired by the firm; we call it an *unsuccessful hire* when the adviser recommends a low-skilled candidate and he is hired by the firm.

Reputation: A fully dynamic model would feature a repeated game in which, at each period, the firm has a new vacant position and the adviser has a new candidate. Players would then use strategies based on what happened in previous periods and their expectations about what might happen in future periods. In this context, we want to understand how the firm can use a relational contract to incentivize the adviser to follow the recommendation standard.[12] To simplify the analysis, we consider a single-period game with the addition of a variable called adviser's *reputation*. Reputation plays a key role to incentivize the adviser, and also serves the roles of summarizing past behavior and predicting future behavior.

We start by assuming that the adviser gains some exogenous payoff from having a higher reputation. This payoff captures the adviser's future benefits from having a higher reputation with the firm (we discuss how and when this payoff arrises endogenously in Section 5). The adviser's reputation goes up after a successful hire and down after an unsuccessful hire. Hence, the adviser's current reputation reflects his past behavior — e.g., an adviser with a longer history of successful hires will have a higher reputation. Reputation also influences his current behavior because the adviser will consider how his lying choices affect his reputation. Consequently, reputation impacts the future value of the relationship between the firm and the adviser (lying choices today affect future reputation and future behavior).

Formally, the adviser starts the period with a reputation $r \in \mathbb{R}_{++} = (0, \infty)$, which is common knowledge, and ends the period with a reputation $\tilde{r} \in \mathbb{R}_{++}$. Following a successful hire, the adviser's reputation goes up to $\tilde{r} = r + g(r)$, where $g(r) > 0$ is the adviser's reputation increase. Following an unsuccessful hire, the adviser's reputation goes down to $\tilde{r} = r - c(r)$, where $c(r) \in (0, r)$ is the adviser's reputation decrease. If the adviser does not send a recommendation or if the firm does not hire the focal candidate, then the adviser's reputation stays at $\tilde{r} = r$. The adviser receives a *reputation payoff* $v(\tilde{r})$ if he ends the period

[12] A relational contract is an informal agreement between players in which, at each period, they choose their actions taking into account how their choices affect the future of their relationship. Players choose to follow this informal agreement when it is on their own interest to do so, even when the informal agreement is not enforceable in court. Therefore, such agreements can be very useful when formal contracts are not feasible. For example, when it is hard (or too costly) for the firm to prove in court that a worker recommended by the adviser has low skill.

with a reputation $\tilde{r}$, where $v(0) = 0$ and $v'(\tilde{r}) > 0$. Therefore, the adviser's reputation payoff is $v(r+g(r))$ after a successful hire, $v(r-c(r))$ after an unsuccessful hire, and $v(r)$ otherwise. We assume that functions g, c, and v are differentiable.

Main Assumption: The central tradeoff for an adviser is his recommendation choice when his candidate has a low skill. If he does not send a recommendation, then his reputation payoff stays at $v(r)$. If he recommends the low-skilled candidate and he is hired by the firm, then the adviser's reputation payoff goes down to $v(r - c(r))$. Therefore, we can define the adviser's payoff loss from an unsuccessful hire as

$$L(r) \;\; \equiv \;\; v(r) - v(r - c(r)). \tag{1}$$

Our main assumption is that this loss is an increasing function of reputation:

Assumption (A1) *The adviser's payoff loss from an unsuccessful hire is an increasing function of his reputation, $L'(r) > 0$.*

Assumption (A1) relates to the notion that reputation is quick to lose. Since advisers with a higher reputation have more to lose, their reputation payoff loss is higher. Throughout the paper, we will always assume that (A1) holds.[13]

Adviser's Bias: The adviser is biased towards his candidate, in the sense that he receives a positive payoff b if his candidate is hired by the firm, and zero otherwise. For example, this bias may come from the level of personal friendship between the adviser and the candidate. This bias b is a random variable with a cumulative distribution function Z. This distribution is continuous, independent of s, with full support in the interval $B = [0, 1]$, and $Z(0) = 0$. The adviser privately knows the realized value of his personal bias b, but the firm only knows its distribution.[14]

Adviser's Utility: The adviser's utility is the sum of his reputation payoff $v(\tilde{r})$ and, if his candidate is hired, the bias payoff b. Therefore, after a successful hire his utility is

[13] We can rewrite condition $L'(r) > 0$ as $c'(r) > 1 - \frac{v'(r)}{v'(r-c(r))}$. Several functional forms of v and c satisfy this condition. In particular, the assumption always holds if v is convex (so that $0 \geq 1 - \frac{v'(r)}{v'(r-c(r))}$) and c is strictly increasing ($c' > 0$); e.g., if $v(r) = r$ and $c(r) = \delta r$ for some $\delta \in (0, 1)$, then $L'(r) = \delta$ and the condition holds. Moreover, c' could be negative (but not too negative) if v is strictly convex. Conversely, v could be concave as long as c' is sufficiently large.

[14] Previous literature has considered the use of referral bonus to incentivize advisers, which can create a bias. Our approach complements this literature in two ways. First, when the firm offers the bonus, the firm knows exactly the bias of the adviser. It is then easier for the firm to infer the adviser's behavior, and it is easier to have a corner solution in which the adviser does not lie. In our setup, the adviser's bias is private information, hence the firm is more concerned about the true motives of the adviser, which we believe is a realistic assumption in many scenarios. Second, the literature has explored how a direct payment can motivate advisers. We complement this literature by explaining how our relational contract can endogenously provide incentives to the adviser.

$v(r+g(r))+b$, and after an unsuccessful hire his utility is $v(r-c(r))+b$. In all other cases (recommending a candidate who is not hired or not recommending a candidate), the adviser keeps his reputation r; in these remaining cases, his utility is then $v(r)+b$ if his candidate is hired and $v(r)$ if his candidate is not hired.

Adviser's Strategy: After privately observing his personal bias b and his candidate's skill s, and given his reputation r, the adviser chooses whether to recommend the candidate. We represent the adviser's strategy[15] as a function $\sigma^A : B \times S \times \mathbb{R}_{++} \rightarrow M$. The adviser chooses the strategy that maximizes his expected utility.

Firm's Payoff and Strategy: If the firm hires the adviser's focal candidate, then it receives a payoff s_H if he has high quality and s_L if he has low quality. If the firm hires the alternative candidate, then it receives a payoff x from this outside option. We can interpret x as the expected skill of the best alternative candidate currently available to the firm. Payoff x is a random variable with cumulative distribution function F. We assume that F is continuously differentiable, with probability density function f, full support in the interval $X \equiv [0,1]$, independent of b, s and m. The firm privately knows the realized value x of the alternative candidate. However, the adviser only knows the distribution F.

After observing the realized quality x of the alternative candidate, the adviser's reputation r and his message m, the firm needs to choose which candidate to hire. We denote this action as $a \in A \equiv \{0,1\}$, where $a=1$ if the firm hires the adviser's candidate, and $a=0$ if the firm hires the outside candidate. We represent the firm's strategy as a function $\sigma^F : X \times \mathbb{R}_{++} \times M \rightarrow A$.[16] The firm chooses the strategy that maximizes its expected payoff.

Timing of the Game:

1. Adviser with reputation r privately observes his personal bias b and the focal candidate's skill s.

2. Adviser sends a message $m \in M$ to the firm.

3. The firm observes message m and the realized value x of the outside option. The firm then chooses which candidate to hire, $a \in A$.

4. The adviser's reputation is adjusted according to the outcome and payoffs are realized.

Equilibrium: Our equilibrium concept is perfect Bayesian equilibrium (PBE). Without loss of generality, if a player is indifferent between multiple optimal choices, we break ties as

[15]We are considering pure strategies without loss of generality. We will show that, in equilibrium, only a measure zero of adviser's types would be willing to mix between recommending or not.

[16]We are considering pure strategies without loss of generality. We will show that, in equilibrium, only a measure zero of firm's types would be willing to mix between the two candidates.

follows: (i) the firm hires the adviser's candidate if it is indifferent between the two candidates, (ii) the adviser does not recommend a low-skilled candidate if it is indifferent between recommending or not. Henceforth, we refer to PBE equilibria featuring these characteristics simply as equilibrium.

3 Equilibrium Characterization

Our first result shows that the adviser and the firm have optimal strategies characterized by thresholds. We provide an intuitive overview before presenting the formal result.

Adviser's Equilibrium Strategy: After observing that his candidate has a high skill, the adviser always sends a recommendation, independently of his bias and reputation. After observing a low skill, the adviser faces a tradeoff: if he lies and sends a recommendation, he increases the chances that his candidate will be hired (and he will receive the bias payoff b), but he will have to pay the reputation payoff loss $L(r)$ if he is hired. Therefore, his behavior is characterized by an equilibrium threshold $b^* \in (0,1]$. If the adviser's bias is high ($b > b^*$), then he lies and sends a recommendation. An adviser with a low bias $b \leq b^*$ prefers to not send a recommendation.

Firm's Belief and Equilibrium Strategy: To characterize the firm's equilibrium strategy, we need to compute its belief after observing a recommendation and after not observing a recommendation. This belief needs to be consistent with the adviser's equilibrium strategy.

After observing no recommendation, the firm becomes certain that the focal candidate has a low skill. Recall that the firm receives payoff s_L if it hires a low-skilled candidate. The firm's expected payoff from hiring the alternative candidate is x. Consequently, the firm will hire the alternative candidate if he provides a higher expected payoff ($x > s_L$), otherwise he will hire the focal candidate ($x \leq s_L$). Variable x is distributed according to F; hence, following no recommendation, the firm hires the focal candidate with probability $F(s_L)$.

After observing a recommendation, the firm must take into account that the adviser might be lying. An adviser with a high-skilled candidate always tells the truth, while an adviser with a low-skilled candidate lies if his bias is above the threshold b^*. Since the adviser's bias follows the distribution Z, an adviser who has a low-skilled candidate lies and sends a recommendation with probability $1 - Z(b^*)$. Let $s_R(b^*)$ be the firm's belief about the focal candidate's expected skill after observing a recommendation, when the adviser uses the bias threshold b^*. Bayes' rule implies that (see the proof of Proposition 1 for details)

$$s_R(b^*) \quad = \quad \frac{ps_H + (1-p)(1-Z(b^*))s_L}{p + (1-p)(1-Z(b^*))}. \tag{2}$$

Note that the firm's belief $s_R(b^*)$ is an increasing function of b^*. If the adviser always lies ($b^* = 0$), then the firm keeps its prior belief $s_R(0) = \hat{s}$ and does not trust a recommendation. That is, the firm views the recommended candidate as an average candidate $\hat{s}$, since the adviser always sends a recommendation. If the adviser never lies ($b^* = 1$), the firm fully trusts a recommendation and it is sure that the candidate has a high skill, $s_R(1) = s_H$. If the adviser lies with positive probability ($b^* < 1$), then the firm's belief s_R will be strictly less than s_H; that is, the firm "discounts" the adviser's recommendation and takes the information with a grain of salt. The firm's expected payoff from the recommended focal candidate is $s_R(b^*)$, while the payoff from the alternative candidate is x; thus, the firm will hire the alternative candidate if $x > s_R(b^*)$, otherwise it will hire the focal candidate. Hence, following a recommendation, the firm hires the focal candidate with probability $F(s_R(b^*))$.

Computing the Equilibrium Threshold: Our previous discussion shows that equilibrium beliefs and strategies can be fully characterized by the threshold b^*. We want to compute the value b^* such that strategies and beliefs are consistent with each other.

Suppose the firm believes that the adviser is using some threshold b^*. Recall that an adviser with a high-skilled candidate always sends a recommendation. Hence, consider the remaining case: an adviser with a low-skilled candidate. If he chooses no recommendation, then: the firm correctly interprets that the candidate has a low skill, the focal candidate is hired with probability $F(s_L)$, and the adviser keeps his current reputation. No recommendation yields an expected utility $v(r)+bF(s_L)$. If the adviser choose to send a recommendation, then the firm updates its belief to $s_R(b^*)$. With probability $F(s_R(b^*))$ the focal candidate is hired, in which case the adviser receives his bias payoff b but his reputation payoff decreases to $v(r - c(r))$. With probability $1 - F(s_R(b^*))$ the focal candidate is not hired, in which case the adviser does not receive the bias payoff and his reputation payoff stays at $v(r)$. The adviser strictly prefers to lie if it yields a strictly higher expected utility:

$$\begin{aligned}
\text{expected utility from lying} &> \text{expected utility from not lying} \\
\iff F(s_R(b^*))\,[b + v(r - c(r))] + (1 - F(s_R(b^*)))v(r) &> bF(s_L) + v(r) \\
\iff b\,[F(s_R(b^*)) - F(s_L)] &> F(s_R(b^*))\,[v(r) - v(r - c(r))] \\
\iff b\,[F(s_R(b^*)) - F(s_L)] &> F(s_R(b^*))L(r) \\
\iff b\left[\frac{F(s_R(b^*)) - F(s_L)}{F(s_R(b^*))}\right] &> L(r). \qquad (3)
\end{aligned}$$

The LHS of (3) captures the benefit from lying. The term $[F(s_R(b^*)) - F(s_L)]\,/F(s_R(b^*))$ captures the relative increase in the probability of the firm hiring the focal candidate when the adviser lies. Note that the LHS is increasing in his personal bias b. This is why the adviser's optimal strategy must take the form of a threshold on his bias: he lies if his bias is high enough.

The RHS of (3) is the reputation payoff loss from an unsuccessful hire, defined by (1), which depends on the adviser's reputation. In the proof of Proposition 1, we define a cutoff $\bar{r}$ which separates two possible cases. For every adviser with a high reputation $r \geq \bar{r}$, the reputation loss $L(r)$ is weakly greater than the maximum benefit from lying, thus we have a corner solution: $b^* = 1$ and this adviser will not lie. For every adviser with a low reputation $r < \bar{r}$, we have an interior solution: there is a unique $b^* \in (0, 1)$ which solves

$$b^* \left[\frac{F(s_R(b^*)) - F(s_L)}{F(s_R(b^*))} \right] \quad = \quad L(r). \tag{4}$$

Therefore, an adviser with a low-skilled candidate and a large bias $b > b^*$ strictly prefers to lie and send a recommendation, while an adviser with a small bias $b \leq b^*$ sends no recommendation. The following proposition summarizes our results.

Proposition 1 *The game has a unique equilibrium.*

(i) The adviser's equilibrium strategy is the following:

- *(a) An adviser with a high-skilled candidate always sends a recommendation.*
- *(b) An adviser with a low-skilled candidate sends a message according to a threshold b^*: he lies and sends a recommendation if he has a high bias ($b > b^*$), and does not send a recommendation if he has a low bias $b \leq b^*$.*

(ii) The firm's equilibrium strategy is the following:

- *(a) After receiving no recommendation, the firm updates its belief to s_L; it hires the focal candidate if $x \leq s_L$ and the alternative candidate otherwise.*
- *(b) After receiving a recommendation, the firm takes into account the lying probability and only updates its belief to $s_R(b^*)$, given by (2); it hires the focal candidate if $s_R(b^*) \geq x$ and the alternative candidate otherwise.*

4 Main Results

In this Section we present our main results.

4.1 Adviser's Reputation

We first consider how the adviser's current reputation r affects equilibrium behavior. The next proposition shows the intuitive result that the adviser lies less if his reputation is higher, which benefits the firm.

Proposition 2 *In equilibrium, an adviser's probability of lying decreases in his reputation r. Consequently, the firm's ex-ante payoff is increasing in the adviser's reputation.*

According to our main assumption (A1), the payoff loss $L(r)$ is an increasing function of reputation. An increase in r increases the RHS of (4), hence we need to also increase b^* to increase the LHS of (4). A higher equilibrium cutoff b^* means that the adviser is less likely to lie. Intuitively, our model captures the idea that an adviser with a higher reputation has "more to lose" when he lies to the firm. Because an adviser with a higher reputation is less likely to lie, he provides a more informative signal[17] to the firm. With more information about the focal candidate, the firm can make a better hiring decision and increase its expected payoff.

4.2 Firm's Outside Market

The firm has to choose between the focal candidate and the outside option x, which follows the distribution F. This distribution can be interpreted as capturing the "strength" of the firm, and describing the competition that the focal candidate faces. Different models might define strength in different ways, depending on their application. For example, some theoretical and empirical models may only focus on the average skill of the distribution, or use concepts such as first-order stochastic dominance to rank different markets. It turns out that, in our model, a crucial factor is the overall shape of the distribution. We focus on two types of markets, defined as follows.

Definition: *We say that the firm has access to a top-heavy (bottom-heavy) outside market if the probability density function f is strictly increasing (decreasing).*

Intuitively, a firm with a top-heavy market has better access to top outside candidates — outside candidates with a high ability x are very likely to want to come to the firm. A firm with a bottom-heavy market has very limited access to these top outside candidates. It is important to note that the shape of f defines the shape of F. A top-heavy market (increasing f) implies a strictly convex F; a bottom-heavy market (decreasing f) implies a strictly concave F.

[17] A signal $S1$ is considered more informative than another signal $S2$ if every Bayesian decision maker prefers $S1$ to $S2$.

We will show that there is an important equilibrium interaction between the firm's market and the adviser's reputation. We start by recalling a cornerstone result in economics. Consider an individual who must choose between two lotteries which have the same expected return. The first lottery is safer: it pays a small winning prize, but it is very likely that the individual will win this prize. The second lottery is riskier: it pays a much larger winning prize, but it is very unlikely that the individual will win this prize. In this case, a risk averse individual prefers the safer lottery over the riskier lottery, while the opposite is true for a risk lover. Next, we argue that a similar logic applies to our model.

To provide some intuition, it is useful to first contrast two extremes: an adviser facing a zero reputation loss and an adviser facing a very high reputation loss. If there is no reputation payoff loss ($L(r) = 0$), then the adviser has nothing to lose. Therefore, when his candidate has a low skill, the adviser always lies and sends a recommendation. The firm always fully discounts this recommendation, keeps its prior $\hat{s} = ps_H + (1-p)s_L$, and hires the candidate with probability $F(ps_H + (1-p)s_L)$. If the adviser faces a very high reputation payoff loss ($L(r) \geq 1$), then he always tells the truth and faces a gamble. With probability p the candidate will be high-skilled and he will send a recommendation; the firm fully trusts the recommendation, updates its belief to s_H, and hires the focal candidate with the higher probability $F(s_H)$. With probability $(1-p)$ the candidate has a low skill and the adviser does not send a recommendation; the firm updates its belief to s_L and hires the focal candidate with the lower probability $F(s_L)$. The focal candidate's hiring probability is strictly higher under this gamble if and only if

$$\begin{aligned} &\text{Gamble Hiring Probability} > \text{No-Gamble Hiring Probability} \\ \iff\ & pF(s_H) + (1-p)F(s_L) \quad > F(ps_H + (1-p)s_L). \end{aligned} \tag{5}$$

We can then directly apply the definition of strict concavity and strict convexity. If F is strictly convex, then inequality (5) holds and the gamble yields a strictly higher hiring probability. If F is strictly concave, then the opposite holds and the gamble yields a strictly lower hiring probability.

In other words, when the adviser truthfully provides information to the firm, he faces a gamble: sometimes he provides a recommendation and the hiring probability goes up; sometimes he provides no recommendation and the hiring probability goes down. When the firm has a top-heavy market (F is strictly convex), it is as if the adviser is risk lover: the gamble increases the overall hiring probability. When the outside market is bottom-heavy (F is strictly concave), it is as if the adviser is risk averse: the gamble decreases the overall hiring probability. This logic extends to any lying probability and it replicates well known results from the persuasion literature.[18]

It is easy to visualize these results. In Figure 1(a), we have a firm with a top-heavy market (F is strictly convex). In the example, $s_H = 0.9$, $s_L = 0.1$, $p = \frac{1}{2}$, and $\hat{s} = \frac{1}{2}$. If the adviser's potential reputation payoff loss is zero, then the firm keeps its prior belief and the hiring probability is $F(\hat{s})$, which is point A in the graph. When the adviser has an intermediate reputation and only lies sometimes, with some probability the adviser sends a recommendation and the hiring probability goes up to point C; with complement probability the adviser sends no recommendation and hiring probability goes down to point B. On average, the hiring probability in this gamble is given by point D, which is higher than the no-gamble probability A because the function F is strictly convex. If the adviser has a very high reputation such that he never lies, then we have two competing effects. When the adviser sends a recommendation, the firm will trust more this recommendation (will have a higher belief about the candidate's expected skill) and will hire the focal candidate with a higher probability — this is point C' in Figure 1(b), which is higher than point C in Figure 1(a). However, because the adviser is less likely to lie, he is less likely to send a recommendation — the probability of reaching point C' is lower than the probability of reaching point C. Since F is strictly convex, the overall hiring probability of the gamble goes up to point D' in Figure 1(b).

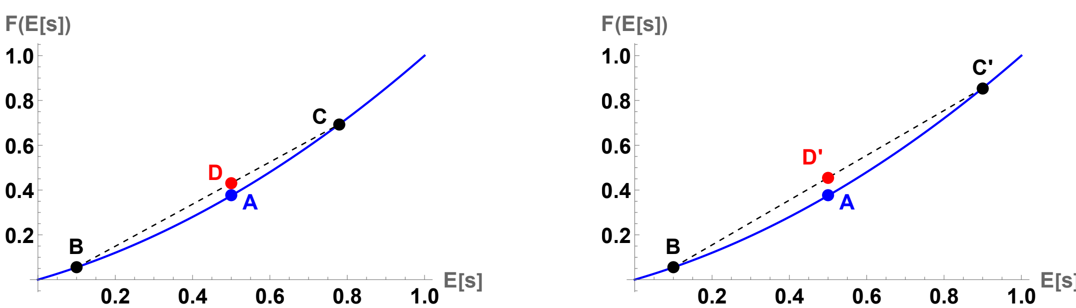

(a) Intermediate Reputation (b) High Reputation

Figure 1: Information as a Gamble, Top-Heavy Firm with $s_H = 0.9$, $s_L = 0.1$, $p = \frac{1}{2}$, $\hat{s} = \frac{1}{2}$.

In Figure 2(a), we have a firm with a bottom-heavy market (F is strictly concave). If the adviser's potential reputation payoff loss is zero, then the firm keeps its prior belief and the hiring probability is given by point A in the graph. When the reputation is intermediate, with some probability the adviser sends a recommendation and hiring probability goes up to point C; with complement probability the adviser sends no recommendation and hiring

[18] In the persuasion literature, if the sender's payoff is a strictly convex (concave) function of the receiver's posterior belief, then the sender wants to provide more (less) information to the receiver — e.g., see Kamenica and Gentzkow (2011). In our setup, the shape of the adviser's payoff function arises naturally from the firm's outside market (the firm's access to alternative candidates).

probability goes down to point B. On average, the hiring probability in this gamble is given by point D, which is lower than the no-gamble probability A because the function F is strictly concave. If the reputation is high and the adviser never lies, then we have two competing effects. When the adviser sends a recommendation, the firm will trust more this recommendation (will have a higher belief about the candidate's expected skill) and will hire the focal candidate with a higher probability — this is point C' in Figure 2(b), which is higher than point C in Figure 2(a). However, because the adviser is less likely to lie, he is less likely to send a recommendation — the probability of reaching point C' is lower than the probability of reaching point C. Since F is strictly concave, the overall hiring probability of the gamble goes down to point D' in Figure 2(b).

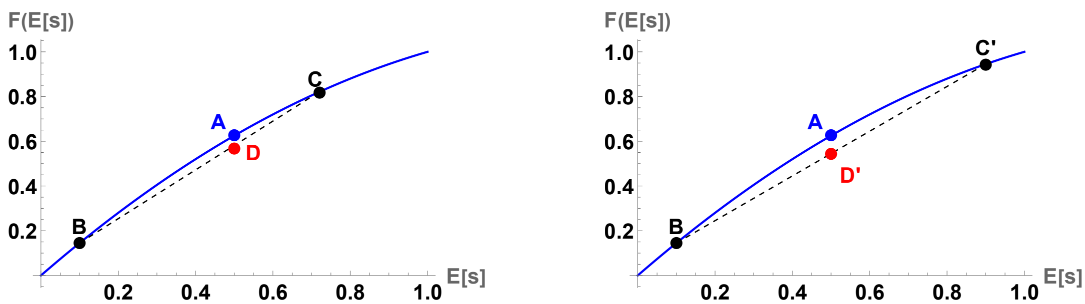

(a) Intermediate Reputation (b) High Reputation

Figure 2: Information as a Gamble, Bottom-Heavy Firm with $s_H = 0.9$, $s_L = 0.1$, $p = \frac{1}{2}$, $\hat{s} = \frac{1}{2}$

It is important to emphasize that these results regarding the effects of reputation on hiring probability hold independently of whether the focal candidate is an underdog or a frontrunner.[19] That is, independently of whether the firm has a prior belief about the skill of the focal candidate that is low or high. The main driving force is the shape of the skill distribution of the alternative candidate, as captured by the definition of top-heavy and bottom-heavy markets. Our results highlight that only looking at the average outside candidate might not be enough to explain the main economic forces driving recommendations and hiring decisions.

Proposition 3 summarizes how the firm's market interacts with the adviser's reputation to create an assortative matching.

Proposition 3 *In equilibrium:*

(i) If the firm has a top-heavy market, then the probability of hiring the adviser's candidate is an increasing function of his reputation r.

[19]See also Alonso and Câmara (2016) for a related result in their model of competition between two political candidates prior to an election.

(ii) If the firm has a bottom-heavy market, then the probability of hiring the adviser's candidate is a decreasing function of his reputation r.

(iii) Consequently, there is assortative matching: a firm with a top-heavy market is more likely to hire a candidate from an adviser with a stronger reputation, while a firm with a bottom-heavy market is more likely to hire a candidate from an adviser with a weaker reputation.

Firms trust more an adviser with a higher reputation because he is less likely to lie — he has more to lose. To be a successful candidate in a firm with a top-heavy market, telling the truth is more important because the competition is more intense (function F is convex). Therefore, strong advisers have better chances with strong firms.

4.3 Low Reputation Trap

Our next result shows that the adviser might face a reputation trap. If he starts the period with a low reputation r, then he expects to end the period with a strictly lower reputation $\tilde{r}$. Conversely, if he starts the period with a high reputation r, then he expects to end the period with a strictly higher reputation $\tilde{r}$.

First, consider an adviser with a high reputation. To guarantee that his reputation is expected to go up, we need this adviser to lie with a sufficiently low probability. A sufficient condition is that the reputation cost of an unsuccessful hire is high, that is, $L(r) \geq 1$ for some $r \in \mathbb{R}_{++}$. In this case, we can guarantee that advisers with a high reputation will lie with a sufficiently low probability and their reputations are expected to go up. This condition complements our motivation that reputation is quick to lose. Hence, advisers with a high reputation are particularly worried about losing their reputation.

Second, consider an adviser with a low reputation. To guarantee that his reputation is expected to go down, a sufficient condition is that reputation is slow to build, that is, the reputation gain $g(r)$ for a successful hire is small. To understand why this condition is relevant, consider the extreme case in which an adviser with a low reputation always recommends his candidate. With probability p the candidate will actually have a high skill and the adviser will (correctly) recommend him; if the firm hires this candidate, then the adviser's reputation will increase by $g(r)$. With probability $1 - p$ the candidate will have low skill and the adviser will lie and recommend him; if the firm hires this candidate, then the adviser's reputation will decrease by $c(r)$. If the gain $g(r)$ is much bigger than the loss $c(r)$, then the expected reputation of the adviser will actually go up, even though he always recommends his low-skilled candidates. Therefore, a sufficient condition is that the gain $g(r)$ is small for advisers with a low reputation.

Proposition 4 *Suppose that the reputation cost of an unsuccessful hire is high and reputation is slow to build, that is, $L(r) \geq 1$ for some $r \in \mathbb{R}_{++}$ and $g(r)$ is sufficiently small for advisers with a low reputation. Then, there exists a reputation cutoff $r_T \in \mathbb{R}_{++}$ such that:*

(i) If an adviser starts the period with a low reputation, then his reputation is expected to go down: $r < r_T$ implies $E[\tilde{r}] < r$.

(ii) If an adviser starts the period with a high reputation, then his reputation is expected to go up: $r > r_T$ implies $E[\tilde{r}] > r$.

A relevant example is the following: $v(r) = r$, $c(r) = \delta r$ for some $\delta \in (0, 1)$, and $g(r) = \kappa r$ for some $\kappa > 0$. In this case, the reputation cost of an unsuccessful hire is high: $L(r) = r - (r - \delta r) = \delta r$, so that $L(r) \geq 1$ if r is high. If reputation is slow to build (if κ is sufficiently small), then there exists a cutoff r_T such that all advisers with a low reputation $r < r_T$ are expected to further lose reputation, while all advisers with a high reputation $r > r_T$ are expected to gain reputation.[20]

4.4 Market Competition

The adviser's focal candidate is competing against the outside candidate, whose skill x is distributed according to F. What happens with equilibrium behavior if we increase the degree of competition, by increasing the likelihood that the outside candidate has a better skill? It is useful to define two special cases. Recall that $\hat{s}$ is the expected skill of the focal candidate. Accordingly, we denote "top-candidates" the outside candidates with skill $x > \hat{s}$, and "bottom-candidates" the candidates with skill $x < \hat{s}$. We say that the firm has access to better top-candidates if we change the outside market from F to $\hat{F}$, so that the distribution of bottom-candidates remains the same, but the distribution of top-candidates improves — the firm is more likely to find top-candidates with higher skill. Similarly, we say that the firm has access to better bottom-candidates if we change the outside market from F to $\hat{F}$, so that the distribution of top-candidates remains the same, but the distribution of bottom-candidates improves — the firm is more likely to find bottom-candidates with higher skill. The next definition formalizes this notion.

Definition: *The firm has access to better top-candidates if we change the outside market from F to $\hat{F}$, where $F(x) = \hat{F}(x)$ for all $x \in [0, \hat{s}]$, and $F(x) > \hat{F}(x)$ for all $x \in (\hat{s}, 1)$. The firm has access to better bottom-candidates if we change the outside market from F to $\hat{F}$, where $F(x) = \hat{F}(x)$ for all $x \in [\hat{s}, 1]$, and $F(x) > \hat{F}(x)$ for all $x \in (0, \hat{s})$.*

[20] In the proof of Proposition 4, we show that the exact condition is $\kappa < \frac{(1-p)\delta}{p}$.

Proposition 5 shows how changes in the outside market affect equilibrium behavior.

Proposition 5 *Suppose that the adviser lies with a probability strictly between zero and one, $b^* \in (0,1)$. Then:*

(i) If the firm has access to better top-candidates, then the adviser lies less and the firm's expected payoff goes up.

(ii) If the firm has access to better bottom-candidates, then the adviser lies more and the firm's expected payoff might go up or down.

To understand the result, consider first an improvement in the distribution of top-candidates and recall equilibrium condition (4). Under the original market distribution F, the benefit from lying is $b^* \left[\frac{F(s_R(b^*))-F(s_L)}{F(s_R(b^*))}\right]$, or simply $b^* \left[1 - \frac{F(s_L)}{F(s_R(b^*))}\right]$. Under the new market distribution $\hat{F}$, if we were to hold b^* constant, the new benefit from lying would be $b^* \left[1 - \frac{F(s_L)}{\hat{F}(s_R(b^*))}\right]$. This benefit is smaller because $\hat{F}(s_R(b^*)) < F(s_R(b^*))$. Because the lying benefit is lower, we need to increase b^* to maintain the equilibrium condition (4). Note that the firm always benefits from better top-candidates: the adviser lies less and the firm can draw better outside options. A double-win for the firm.

Consider now an improvement in the distribution of bottom-candidates. Under the new market distribution $\hat{F}$, if we were to hold b^* constant, the new benefit from lying would be $b^* \left[1 - \frac{\hat{F}(s_L)}{F(s_R(b^*))}\right]$. This benefit is bigger because $\hat{F}(s_L) < F(s_L)$. Because the lying benefit is higher, we need to decrease b^* to maintain the equilibrium condition (4). Interestingly, the firm faces a tradeoff when it has access to better bottom-candidates. More competition on the bottom creates a direct benefit (the firm can draw better outside options) but an indirect cost (the adviser lies more). Which effect dominates will depend on the parameters of the model. In the proof of the proposition, we provide an example in which the firm's payoff goes down with an improvement in the outside market. Loosely speaking, the negative effect dominates if the change in F is sufficiently concentrated around s_L.

Our results imply a "rich get richer" scenario: if a firm is able to attract more top-quality candidates, then this firm will also receive more informative advice and will be able to make better hiring decisions. Moreover, less attractive firms might get stuck in transition. Suppose that when a less attractive firm starts to become more attractive, it first attracts better bottom-quality alternative candidates. This firm will then suffer the drawback of receiving less reliable recommendations, compromising its hiring decisions and curbing its transition to become a very attractive firm.

The above results show how improvements in the quality of the alternative candidate affect equilibrium behavior. We now ask the question: holding constant the expected quality of the alternative candidate, what happens if the firm has access to more reliable information about the quality of the alternative candidate? We represent more information with a mean-preserving spread of the skill distribution: the firm becomes more likely to observe lower and higher values of skill x. Formally, we adapt the rotation approach of Johnson and Myatt (2006):

Definition: *Outside market $\hat{F}$ provides to the firm more reliable information about the alternative candidate than outside market F if $F(x) < \hat{F}(x)$ for all $x \in (0, \hat{s})$, $F(x) > \hat{F}(x)$ for all $x \in (\hat{s}, 1)$, and the expectation of x is the same in both markets.*

In other words, outside market $\hat{F}$ is a mean-preserving rotation of F at $\hat{s}$ — see the example in Figure 3(a), which rotates F at $\hat{s} = 0.5$. The rotation moves density away from $\hat{s}$ and towards the top and bottom of the skill distribution — see Figure 3(b).

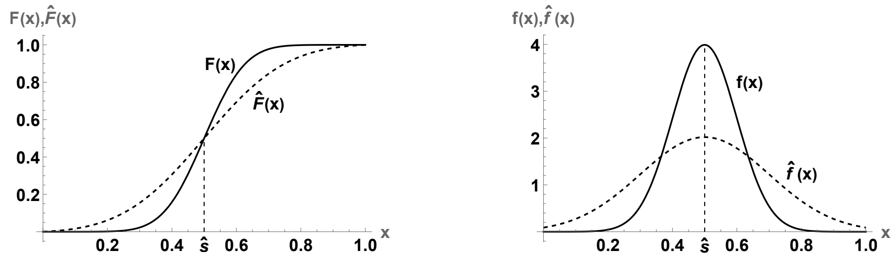

(a) Rotation of F (b) Spread of f

Figure 3: Example of More Reliable Information about the Alternative Candidate

Proposition 6 *Consider markets with interior solutions. If the firm has more reliable information about the alternative candidate, then the adviser will lie less about the focal candidate.*

Proposition 6 shows that equilibrium behavior can result in self-fulfilling expectations. If advisers believe that a firm will receive more reliable information about alternative candidates, then advisers lie less and provide more reliable information about their focal candidates. Together Propositions 5 and 6 show how strategic communication can amplify and sustain differences across firms. Some firms may have easier access to many top-candidates and receive more reliable recommendations, while other firms might be relegated to bottom-candidates and unreliable recommendations.

5 Extensions of the model

In this section, we discuss some possible extensions of our model.

5.1 Endogenous value of reputation

In our benchmark model, we consider a single period model and assume that the adviser receives an exogenous benefit from ending the period with a higher reputation. It is beyond the scope of this paper to formally solve a fully dynamic model in which, at each period, the firm has a new vacant position and the adviser has a new candidate. However, we can discuss how the value of reputation may arise endogenously in such scenario.

First, consider the case of a firm that has a top-heavy outside market. Proposition 2 implies that if keeping a higher reputation is valuable, then an adviser with a higher reputation will lie less. Moreover, Proposition 3(i) shows that, for a top-heavy firm, the probability of hiring a focal candidate is an increasing function of the adviser's reputation. Therefore, at each period, a higher reputation translates to a higher hiring probability and, hence, a higher bias payoff for the adviser. Consequently, reputation endogenously generates a benefit for the adviser, in the form of a higher hiring probability in the future. Reputation thus provides a complete self-contained bootstrap mechanism: a higher reputation is valuable because it induces the adviser to lie less, which increases future hiring probabilities; hence, advisers choose to lie less to nurture their valuable reputation.

Now consider a firm with a bottom-heavy outside market. In this case, the probability that the firm hires the focal candidate is a decreasing function of the adviser's reputation; i.e., reputation has a negative effect on future hiring probabilities. Because reputation decreases hiring probabilities, firms with bottom-heavy outside markets would need to provide some form of an additional incentive to advisers with high reputations, to counterbalance the negative effect of reputation on hiring probabilities. For example, if the adviser is an employee of the firm, then the firm could implement a policy in which advisers with a higher reputation are more likely to receive pay raises and promotions. It is important to note that an adviser with a higher reputation is more valuable to the firm because he is more trustworthy on his recommendations — Proposition 2. Therefore, the firm has an endogenous incentive to keep and nurture an adviser with a higher reputation.

In summary, the endogenous value of reputation can generate another "rich get richer" effect. Firms with access to top-heavy outside markets are in a privileged position, as reputation endogenously has a positive value via increased hiring probability. Hence, such firms may be able to avoid having to use other incentives to motivate advisers. In contrast, firms with a bottom-heavy market might need to provide additional incentives to advisers with a higher reputation, to compensate for the fact that a higher reputation leads to a lower hiring probability.

5.2 Firm does not know if the candidate has an adviser

In our benchmark model, it is common knowledge that the focal candidate has an adviser. In this case, the firm correctly updates its belief to s_L whenever there is no recommendation, since the adviser only chooses to not send a recommendation if the candidate has a low skill.

In some scenarios, the firm might not know if the focal candidate has an adviser or not. For concreteness, suppose that, at the beginning of the period, with probability q the focal candidate has an adviser and with probability $1-q$ he does not have an adviser. The firm knows the probability q, but it cannot directly observe if the candidate has an adviser or not. In this case, the firm must be more careful when evaluating a candidate without a recommendation. No recommendation might be the result of simply not having an adviser (in which case the firm should keep its prior belief $\hat{s}$), or it might be the result of the adviser choosing not to recommend the candidate (in which case the firm's belief should go down to s_L). Let $s_{NR}(b^*)$ be the firm's belief after observing no recommendation.[21] In equation (4), we change s_L to $s_{NR}(b^*)$, and all our qualitative results continue to hold.

5.3 Noisy Signal

In our benchmark model, we assume that the adviser perfectly observes the skill $s \in \{s_L, s_H\}$ of the focal candidate. We can extend the model by assuming that the adviser only observes a noisy signal of this skill, with precision $\alpha \in (0.5, 1]$. Formally, suppose that the adviser observes a binary signal $\omega \in \{L, H\}$, with $\Pr(\omega = H|s_H) = \Pr(\omega = L|s_L) = \alpha$ and $\Pr(\omega = H|s_L) = \Pr(\omega = L|s_H) = 1-\alpha$. A perfectly precise signal $\alpha = 1$ means that the adviser perfectly observes the candidate's skill. A lower precision means that the adviser sometimes makes an honest mistake and may incorrectly evaluate the skill of a candidate (he might observe signal H when the candidate has a low skill or observe signal L when the candidate has a high skill).

We continue to assume that the firm can observe the focal candidate's skill on-the-job after hiring him, and the adviser gains reputation after a successful hire and loses reputation after an unsuccessful hire. However, the firm cannot observe the adviser's noisy signal realization ω. Therefore, when the precision is not one, the adviser might observe a high signal H and correctly recommend the candidate, but the candidate might turn out to be an unsuccessful hire and result in a reputation loss.

We can use Bayes' rule to adjust equation (4) and compute the equilibrium of the model. All our qualitative results continue to hold as long as the signal is sufficiently precise. If

[21] Bayesian updating implies that $s_{NR}(b^*) = \frac{q(1-p)Z(b^*)s_L + (1-q)\hat{s}}{q(1-p)Z(b^*) + (1-q)}$. Note that when $q = 1$ (when the focal candidate always has an adviser, as in our benchmark model), this belief is simply s_L.

the signal is too noisy and the adviser makes too many mistakes evaluating the candidate (when α is too close to 0.5), then some equilibrium features start to change. In particular, an adviser with a low bias might refrain from recommending a candidate after observing a high signal H. If this signal is very imprecise, there is a significant chance that the candidate will turn out to be an unsuccessful hire and the adviser will lose reputation, despite the fact that it was a honest mistake. Hence, the adviser will not recommend the candidate after observing a high signal if his bias is too low.

6 Discussion

In this section, we discuss some features and implications of our model.

6.1 Related Applications

Although we motivate our model in the context of labor markets, a situation where it is highly relevant, the theoretical insights from our model broadly apply to other organizational situations where a decision maker relies on the information provided by a biased adviser — for example, markets for technology (Luo 2014), venture capital investments (Gompers and Lerner 2001), or corporate boards (Hermalin and Weisbach 1998). Moreover, our insights could also be applied to other market scenarios in which an intermediary recommends a product or service to a potential buyer. For instance, online sale platforms, real-estate brokers, and product reviewers in social media.

For example, we could adapt our model to consider a social media influencer (adviser) who can recommend a focal product to a potential buyer (someone following the influencer on social media). The potential buyer is the decision maker in this case — she will choose whether to buy the focal product or an alternative product. We can interpret the influencer's bias as a measure of how much the firm producing the focal product is willing to pay the influencer for a successful recommendation (successful marketing campaign). The potential buyer knows that the influencer is biased, but she is not sure about the degree of the bias. The influencer receives a sample of the product and tests it (learns if it is high quality or low quality). Then, the influencer can recommend the product or not. If the influencer's reputation follows our "slow to build but quick to lose" assumption, then influencers with a higher reputation are less likely to recommend a low-quality product. Moreover, if the buyers' outside option is top-heavy (the buyer is very likely to find a high-quality alternative in the market), then the probability of a successful sale increases with the adviser's reputation. Hence, influencers with a high reputation are very effective marketing tools in top-heavy markets, but less valuable in bottom-heavy markets. This rough example shows how our model can serve as a starting point for future work on related applications.[22]

6.2 Absolute versus Relative Quality

Our model highlights the relevance of considering the candidate's absolute skill (quality) level, as well as the candidate's skill relative to the firm's pool of alternative candidates. In our model, the adviser privately observes the candidate's quality s. For example, advisers potentially observe how well the candidate performs related tasks and jobs, his work and study habits, or his performance in a series of tests. The firm then observes the skill x of an outside candidate, and wants to hire the candidate with the highest skill. Therefore, from the firm's point of view, it is important to learn the focal candidate's absolute skill s to make the best hiring decision.

From the adviser's point of view, since he does not observe the realized skill x of the outside candidate, it is important to take into account the focal candidate's skill relative to the distribution of the outside market. A focal candidate with skill s is better than the outside candidate with probability $F(s)$. Hence, $F(s)$ can be viewed as the candidate's skill relative to the market – for example, a focal candidate with relative skill $F(s) = 0.4$ has a 40% chance of being better than the competing candidate. For the adviser's lying decision, the relevant skill is the relative skill of his candidate, as captured by equation (4).

This means that the same adviser would behave very differently when providing recommendations to different firms; although the absolute skill of his focal candidate has a fixed distribution, the relative skill of his candidate can be quite different across firms. To provide some intuition, suppose $s_L = 0.2$ and $s_H = 0.8$. Suppose firm A has an outside market such that $F_A(s_L) = 0.1$ and $F_A(s_H) = 0.7$, while firm B has an outside market such that $F_B(s_L) = 0.1$ and $F_B(s_H) = 0.6$. The only change from firm A to B is that firm B has access to better top candidates. From Proposition 5(i), we know that the adviser will lie less to firm B. Although we are not changing the absolute skill of the focal candidate, the adviser does change its recommendation strategy to take into account the candidate's skill relative to the market of each firm. Conversely, the firm also takes into account that advisers with access to different candidates (candidates with different absolute skills) behave differently, because they have different relatives skills and, consequently, imply different hiring probabilities.

[22]See Pei and Mayzlin (2022) for a recent theoretical model on the topic of social media influencers. Our main departures from their model are that we allow the influencer to first observe the quality of the product and then choose whether to lie, we consider a buyer with private information about the outside option, and we explicitly consider the role of reputation. One important goal of their paper is to understand the firm's strategic offer to the influencer, which is not captured by our benchmark model (we assume an exogenous bias). We leave this promising extension for future work.

A natural question in then why we cannot focus solely on the relative skill in our model. The reason is that we still need to translate the relative skill into an absolute skill in order to compute the expected skill of the focal candidate. For example, suppose that a low relative skill means $F(s_L) = 0.2$ and a high relative skill means $F(s_H) = 0.8$. Suppose that, in equilibrium, whenever the adviser sends a recommendation, there is a 50/50 chance that the candidate has a high or a low skill. The firm wants to hire the focal candidate if his expected skill is greater than the realized outside option x. What is the expected skill of a recommended candidate? Note that this expected skill will be 0.5 if F is a uniform distribution, but it will be a different number if F is strictly concave or strictly convex. Hence, it is relevant to consider both the absolute and the relative skill of candidates.

6.3 Empirical Relevance

Existing empirical work on labor market recommendations has primarily been conducted by studying employee referral programs within firms, programs where recommendations for potential new hires are solicited from a firm's existing workforce (see Schlachter and Pieper, 2019, for a review). Within this literature, scholars have generally focused on understanding three aspects of the referral process: (1) factors which incentivize workers to make referrals, (2) the likelihood referred candidates are hired, and (3) whether referred hires outperform non-referred hires on a variety of performance metrics (e.g., productivity, wage slope, turnover propensity, etc.). With respect to incentivizing workers to refer job candidates, existing literature has primarily focused on referral bonuses and generally finds that bonuses motivate employees to make referrals (e.g., Friebel et al., 2019; Pieper et al., 2017).[23] However, several empirical studies demonstrate that the quality of referrals decreases as the size of the bonus increases (Beaman and Magruder, 2012; Bond et al., 2018; Friebel et al., 2019). A second common finding in existing empirical literature is that referred candidates are more likely to be hired (Burks et al., 2015; Brown et al., 2016; Pallais and Sands, 2016). A third common finding is that referred hires tend to outperform non-referred hires (Burks et al., 2015; Brown et al., 2016; Pieper, 2015), suggesting better matches with the firm (Topa, 2011).

The results of our model are consistent with the second and third common empirical findings. In the equilibrium of our model, the focal candidate is more likely to be hired if he receives a recommendation from the adviser, and a focal candidate hired with a recommen-

[23]Bonuses can take the form of fixed payments or can be contingent payments typically based on a referred hire remaining within the firm for a pre-determined period of time (Beaman and Magruder, 2012). As described by Ekinci (2016: 703), in practice, few firms tie referral bonuses to referral performance, beyond remaining within the firm.

dation has, on average, a higher skill than a focal candidate hired without a recommendation. The first common finding in existing literature is about the use of direct payments as explicit incentives to advisers. It is natural that the empirical literature has focused on explicit contracts between firms and advisers since data on these are more readily available and easier to measure. Our main contribution is then to study the theoretical relevance and consequences of using reputation as a relational contract between firms and advisers. It is important to consider and measure all incentives provided by firms to advisers, to better understand the relative importance and consequences of explicit and implicit incentives.

It may be difficult to measure empirically the three main ingredients of our model – reputation, lies, and the type of the firm (top-heavy or bottom-heavy). However, we provide strong theoretical reasons for empirical researcher to venture into the path of studying relational contracts.[24] Our model points to a clear theoretical direction to search for new data and measures and demonstrates the value of doing so by generating several novel predictions for empirical specialists to consider.

A first set of predictions relate to the lying behavior of advisers, the expected quality of their recommended candidates, and the dynamics of reputation. In scenarios in which "reputation is slow to build but quick to lose," Proposition 2 implies that advisers with a higher reputation will lie less and, consequently, their recommended candidates with have a higher expected quality. Consequently, the ex-post probability of hiring a recommended candidate is an increasing function of the adviser's reputation. Moreover, Proposition 4 shows that advisers with a high reputation tend to further gain reputation, while advisers with a low reputation are stuck in a reputation trap.

A second set of predictions relate to the assortative matching between firms and advisers. Although the ex-post probability of hiring a recommended candidate is an increasing function of the adviser's reputation, the overall ex-ante probability depends on the firm's type. Because advisers with a higher reputation lie less, Proposition 3 shows that, ex-ante, the probability of hiring the adviser's focal candidates is an increasing (decreasing) function of his reputation if the firm has a top-heavy (bottom-heavy) outside market. Whereas the firm's access to the outside labor market has been largely ignored in existing literature, our model highlights how such markets critically shape the supply of truthful referrals, thus calling for closer examination of this factor in empirical studies.

Our third set of predictions explain what happens to lying and hiring behaviors when the firm faces certain changes. According to Proposition 5, if a firm gains better access to top

[24] Similar issue arises when studying explicit and implicit contracts between firms. Lafontaine and Slade (2013) review the empirical literature on contracts between firms. They argue that most of the empirical literature studies explicit contracts because researchers typically analyze data on what is actually written on the contract, but implicit contracts are a common and important feature of modern economic systems.

candidates, then advisers will lie less, which implies that their recommended candidates will have a higher expected quality.[25] If a firm gains better access to bottom candidates, then the result is ambiguous (it can increase or decrease lies). Proposition 6 shows that if the firm has more reliable information about alternative candidates, then advisers will lie less, which implies that their recommended candidates will have a higher expected quality. For example, firms which use a wider variety of personnel selection methods for job candidates – such as pre-hire mental ability and/or personality tests – would become more informed about potential outside job candidates (Schmidt and Hunter, 1998). Thus, our model suggests that empirical researchers should consider how the quality of employee referrals may vary with the joint use of other selection methods.

Our last set of predictions, as explained in Section 5.1, suggest that for firms with a top-heavy market, reputation has an endogenous value because, in this case, a higher reputation leads to a higher hiring probability and a greater bias payoff. This is not the case for bottom-heavy firms; they may need to provide compensation for advisers with a higher reputation. Hence, it should be easier for top-heavy firms to rely on implicit relational contracts instead of explicit payments, while the opposite is true for bottom-heavy firms.

Finally, our model, which focuses on using reputation as a relational contract, enables researchers to extend beyond the empirical context of employee referral systems. It can be the starting point for other models involving recommendations. Firms often solicit and receive recommendations from actors outside the firm, as would often be the case with letters of recommendation (e.g., Abel et al. 2020), formal and informal candidate endorsements (Castilla and Rissing, 2018), or when firms utilize labor market intermediaries (Autor, 2009). These and the other applications discussed in Section 6.1, such as social media influencers, might provide richer and more accessible data to test our model.

6.4 Future Research

In our model, the outside market competition is represented by a probability distribution function over the quality of the outside candidate. This simplification allows us to provide valuable intuition and sharp comparative static results on the effects of different degrees of competition. Our model can be extended to more explicitly address competition. For example, we have considered one focal candidate referred by one adviser to the firm, whereas it is possible that one adviser can refer multiple candidates (e.g., academic adviser referring multiple students in the same year), and the hiring firm may be considering referrals made

[25]Empirical research could look at, for example, changes in the location of the firm's headquarters (or the headquarters of competitors) that could attract or repel candidates of certain skills as proxies for a market change.

by multiple advisers on different candidates. We leave these extensions for future research.

For simplicity, we have assumed that the adviser privately observes an exogenous signal about the quality of the focal candidate, then he strategically chooses whether to lie about this signal. Future research could investigate how the market changes when the adviser first endogenously choose how much information to collect about the candidate, then strategically decides whether to lie about the information collected. This would then become a persuasion game with endogenous information acquisition and costly lies, which is harder to solve. See Alonso and Câmara (2021) and Guo and Shmaya (2021) for recent developments in this area.

Future research may also extend our model, which focused on hiring decisions, to explore how firms can retain employees once an employee is hired. For example, the model of Gambardella, Giarratana and Panico (2010) demonstrates that information asymmetries between an employee and employer with respect to the employee's true skills can create frictions to employee retention. Such spreads in information are plausibly smaller when firms receive truthful recommendations, as is often the case for higher status and established firms, who also presumably have a greater array of complementary assets available to better incentivize high quality employees to stay after hire (Gambardella et al., 2010). In this sense, better quality information ex-ante hiring may be associated with a higher probability of retaining the employee ex-post. Exploring such interplays between selection and retention in more detail may yield additional and nuanced insights that inform labor market strategy and the sustainability of human capital-based competitive advantages.

Scholars may also build on our model to study the optimal design of soliciting and accepting referrals which incorporates the corresponding tradeoffs between the costs associated with providing referrals and the benefits associated with successful hires, and how the optimal balance of costs and benefits may vary with firm characteristics. While the focus of our paper has been on referrals driven by social connections, the implications of our model extend to other situations where firms use intermediaries such as professional recruiters to screen potential job candidates (Cowgill and Perkowski, 2020). Of course, one key difference between screening job candidates by soliciting referrals from current employees and soliciting referrals from third parties such as recruiters, is that in the former scenario there is a possibility of a "referral treatment effect" where referred hires benefit from having a natural mentor in the referrer, which, in turn, may result in better socialization into the job and hence better worker productivity, even if the referred worker is not ex-ante of higher quality (Burks et al., 2015). Thus, while the general insights from our model apply to both referral and intermediary recommendations, when firms recruit from referrals made by current employees they may wish to provide incentives which encourage more recommendations than the incentives associated with referrals made by third parties, as the former provides the potential for the referral treatment effect which may offset ex-ante lower quality workers.

7 Conclusion

We develop a formal model to address the questions of when advisers choose to provide truthful information, and how hiring firms use this information in making hiring decisions. Our model features an adviser who (i) has private information about the quality of a focal job market candidate, (ii) can costly lie about this quality, and (iii) has private information about his own preference bias; moreover, (iv) the firm making the hiring decision has private information about an alternative candidate. By incorporating these four key features into a unified model, we shed light on the intricate interplay between these forces in shaping the strategies and performance of advisers and hiring firms.

We show that adviser's lying behavior and firm's hiring decisions depend on the interplay between the adviser's reputation (and, consequently, his lying cost) and the hiring firm's strength of accessing the outside market. This mechanism creates assortative matching between hiring firms and advisers: stronger hiring firms are more likely to hire candidates referred to them by advisers with stronger reputation, whereas weaker hiring firms are more likely to hire from advisers with lower reputation. Therefore, the access to labor market ought to be a key consideration in assessing job market referrals. Moreover, we also uncover a "reputation trap" for advisers: advisers with a higher reputation end up with even higher reputation, whereas advisers with a lower reputation suffer from further decrease of reputation.

Finally, the strategic choices of advisers can amplify and sustain differences across firms, thus contributing to a "rich get richer" effect. Advisers lie less to firms that have access to better top-candidates and to firms that have better information about alternative candidates, further benefiting these firms. Therefore, we uncover a mechanism of self-fulfilling expectations: if advisers expect the market to provide more reliable recommendations to the firm, then advisers will provide more reliable information to the firm (they will be less likely to lie).

References

[1] Abel, M., R. Burger, P. Piraino. 2020. The value of reference letters: Experimental Evidence from South Africa. *American Economic Journal: Applied Economics* 12(3) 40-71.

5

[2] Alonso, R. and O. Câmara. 2016. Political Disagreement and Information in Elections. *Games and Economic Behavior*, 100: 390-412.

[3] Alonso, R. and O. Câmara. 2021. Organizing Data Analytics. *USC Working Paper.*

[4] Autor, D. 2009. Studies of labor market intermediation. *University of Chicago Press.*

[5] Banks, J.S. 1990. A model of electoral competition with incomplete information. *Journal of Economic Theory* 50(2) 309-325.

[6] Barney, J.B. 1986. Strategic factor markets: Expectations, luck, and business strategy. *Management Science* 32(10) 1231-1241.

[7] Bar-Isaac, H., S. Tadelis. 2008. Seller Reputation. *Foundations and Trends in Microeconomics* 4(4) 273-351.

[8] Beaman, L., J. Magruder. 2012. Who gets the job referral? Evidence from a social networks experiment. *American Economic Review* 102(7) 3574-3593.

[9] Bidwell, M. 2011. Paying more to get less: The effects of external hiring versus internal mobility. *Administrative Science Quarterly* 56(3) 369-407.

[10] Bidwell, M., S. Won, R. Barbulescu, E. Mollick. 2015. I used to work at Goldman Sachs! How firms benefit from organizational status in the market for human capital. *Strategic Management Journal* 36(8) 1164-1173.

[11] Bond, B. M., T. Labuzova, R. M. Fernandez. 2018. At the expense of quality. *Sociological Science* 5 380-401.

[12] Brown, M., E. Setren, G. Topa. 2016. Do informal referrals lead to better matches? Evidence from a firm?s employee referral system. *Journal of Labor Economics* 34(1) 161-209.

[13] Burks, S.V., B. Cowgill, M. Hoffman, M. Housman. 2015. The value of hiring through employee referrals. *The Quarterly Journal of Economics* 130(2) 805-839.

[14] Callander, S., S. Wilkie. 2007. Lies, damned lies, and political campaigns. *Games and Economic Behavior* 60(2) 262-286.

[15] Carmichael, H.L. 1988. Incentives in academics: why is there tenure? *Journal of Political Economy* 96(3) 453-472.

[16] Castilla, E.J. 2005. Social networks and employee performance in a call center. *American journal of sociology* 110(5) 1243-1283.

[17] Castilla, E.J., B.A. Rissing. 2018. Best in class: The returns on application endorsements in higher education. *Administrative Science Quarterly* 64(1) 230-270.

[18] Chung, W., R. Harbaugh. 2019. Biased recommendations from biased and unbiased experts. *Journal of Economics & Management Strategy* 28(3) 520-540.

[19] Cowgill, B., P. Perkowski. 2020. Delegation in Hiring: Evidence from a Two-Sided Audit. *Columbia Business School Research Paper* No. 898

[20] Crawford, V.P., J. Sobel. 1982. Strategic information transmission. *Econometrica* 50(6) 1431-1451.

[21] Dariel, A., A. Riedl, S. Siegenthaler. 2021. Referral hiring and wage formation in a market with adverse selection. *Games and Economic Behavior* 130: 109-130.

[22] Dustmann, C., A. Glitz, U. Schönberg, H. Brücker. 2016. Referral-based Job Search Networks, *The Review of Economic Studies* 83(2) 514-546.

[23] Ekinci, E. 2016. Employee referrals as a screening device. *The RAND Journal of Economics* 47(3) 688-708.

[24] Fafchamps, M., A. Moradi. 2015. Referral and job performance: Evidence from the Ghana colonial army. *Economic development and cultural change* 63(4): 715-751.

[25] Fernandez, R.M., E.J. Castilla, P. Moore. 2000. Social capital at work: Networks and employment at a phone center. *American Journal of Sociology* 105(5) 1288-1356.

[26] Fernandez, R.M., N. Weinberg. 1997. Sifting and sorting: Personal contacts and hiring in a retail bank. *American Sociological Review* 62(6) 883-902.

[27] Friebel, G., M. Heinz, M. Hoffman, N. Zubanov. 2019. What do employee referral programs do? A firm-level randomized controlled trial. *National Bureau of Economic Research.*

[28] Friebel, G., M. Raith. 2004. Abuse of authority and hierarchical communication. *RAND Journal of Economics* 35(2) 224-244.

[29] Gambardella, A., C. Panico, G. Valentini. 2015. Strategic incentives to human capital. *Strategic Management Journal* 36(1) 37-52.

[30] Gambardella, A., M.S. Giarratana, C. Panico. 2010. How and when should companies retain their human capital? Contracts, incentives and human resource implications. *Industrial and Corporate Change* 19(1) 1-24.

[31] Gompers, P., J. Lerner. 2001. The venture capital revolution. *Journal of Economic Perspectives* 15(2) 145-168.

[32] Greenwald, B.C. 1986. Adverse selection in the labour market. *The Review of Economic Studies* 53(3) 325-347.

[33] Groysberg, B., L. Lee, A. Nanda. 2008. Can they take it with them? The portability of star knowledge workers' performance. *Management Science* 54(7) 1213-1230.

[34] Guo, Y., E. Shmaya. 2021. Costly miscalibration. *Theoretical Economics* 16(2), 477-506.

[35] Haltiwanger, J.C., J.I. Lane, J. Spletzer. 1999. Productivity differences across employers: The roles of employer size, age, and human capital. *American Economic Review* 89(2) 94-98.

[36] Hatch, N.W., J.H. Dyer. 2004. Human capital and learning as a source of sustainable competitive advantage. *Strategic Management Journal* 25(12) 1155-1178.

[37] Heath, R. 2018. Why do firms hire using referrals? Evidence from Bangladeshi garment factories. *Journal of Political Economy* 126(4): 1691-1746.

[38] Hermalin, B.E., M.S. Weisbach. 1998. Endogenously chosen boards of directors and their monitoring of the CEO. *American Economic Review* 88(1) 96-118.

[39] Hodler, R., S. Loertscher, D. Rohner. 2014. Persuasion, binary choice, and the costs of dishonesty. *Economics Letters* 124(2) 195-198.

[40] Ishida, J., T. Shimizu. 2019. Cheap talk when the receiver has uncertain information sources. *Economic Theory* 68(2) 303-334.

[41] Johnson, J. P., D. P. Myatt. 2006. On the Simple Economics of Advertising, Marketing, and Product Design. *American Economic Review* 96 (3): 756-784.

[42] Jovanovic, B. 1979. Job matching and the theory of turnover. *Journal of Political Economy* 87(5) 972-990.

[43] Kamenica, E., M., Gentzkow. 2011. Bayesian Persuasion. *American Economic Review* 101(6): 2590-2615.

[44] Kartik, N. 2009. Strategic communication with lying costs. *The Review of Economic Studies* 76(4) 1359-1395.

[45] Kartik, N., M. Ottaviani, F. Squintani. 2007. Credulity, lies, and costly talk. *Journal of Economic Theory* 134(1) 93-116.

[46] Kreps, D.M., J. Sobel. 1994. Signalling. *Handbook of game theory with economic applications 2*: 849-867.

[47] Lafontaine, F., M. E. Slade. 2013. Inter-firm contracts. In Gibbons, R., J. Roberts, editors. *The Handbook of Organizational Economics*. Princeton University Press, 2013.

[48] Lowenstein, R. 1995. Buffett: The Making of an American Capitalist. *Random House Publishing*. New York.

[49] Luo, H. 2014. When to sell your idea: Theory and evidence from the movie industry. *Management Science* 60(12) 3067-3086.

[50] Mailath, G. J., L. Samuelson. 2006. Repeated games and reputations: long-run relationships. *Oxford University Press*.

[51] Malcomson, J. M. 2013. Relational Incentive Contracts. In Gibbons, R., J. Roberts, editors. *The Handbook of Organizational Economics*. Princeton University Press, 2013.

[52] Montgomery, J.D. 1991. Social networks and labor-market outcomes: Toward an economic analysis. *The American Economic Review* 81(5) 1408-1418.

[53] Oyer, P., S. Schaefer. 2010. Personnel economics: hiring and incentives. *National Bureau of Economic Research.*

[54] Pallais, A., E. G. Sands. 2016. Why the referential treatment? Evidence from field experiments on referrals. *Journal of Political Economy* 124(6): 1793-1828.

[55] Pei, A, D. Mayzlin. 2022. Influencing Social Media Influencers Through Affiliation. *Marketing Science* 41(3), 593-615.

[56] Pieper, J.R. 2015. Uncovering the nuances of referral hiring: How referrer characteristics affect referral hires' performance and likelihood of voluntary turnover. *Personnel Psychology* 68(4) 811-858.

[57] Pieper, J. R., J. M. Greenwald, S. D. Schlachter. 2017. Motivating employee referrals: The interactive effects of the referral bonus, perceived risk in referring, and affective commitment. *Human Resource Management* 57(5) 1159-1174.

[58] Rider, C.I., D. Tan. 2014. Labor market advantages of organizational status: A study of lateral partner hiring by large US law firms. *Organization Science* 26(2) 311-631.

[59] Saloner, G. 1985. Old boy networks as screening mechanisms. *Journal of Labor Economics* 3(3) 255-267.

[60] Schlachter, S.D., J.R. Pieper. 2019. Employee referral hiring in organizations: An integrative conceptual review, model, and agenda for future research. *Journal of Applied Psychology* 104(11) 1325-1346.

[61] Schmidt, F. L., J. E. Hunter. 1998. The validity and utility of selection methods in personnel psychology: Practical and theoretical implications of 85 years of research findings. *Psychological bulletin* 124(2) 262.

[62] Seidmann, D.J. 1990. Effective cheap talk with conflicting interests. *Journal of Economic Theory* 50(2) 445-458.

[63] Singh, J., A. Agrawal. 2011. Recruiting for ideas: How firms exploit the prior inventions of new hires. *Management Science* 57(1) 129-150.

[64] Spence, M.A. 1973. Job market signaling. *The Quarterly Journal of Economics* 87(3) 355-374.

[65] Starr, E., M. Ganco, B. Campbell. 2018. Strategic human capital management in the context of cross-industry and within-industry mobility frictions. *Strategic Management*

Journal 39(8) 2226-2254.

[66] Stigler, G.J. 1961. The economics of information. *Journal of Political Economy* 69(3) 213-225.

[67] Stigler, G.J. 1962. Information in the labor market. *Journal of Political Economy* 70(5, Part 2) 94-105.

[68] Topa, G. 2011. Labor markets and referrals Handbook of social economics. *Elsevier,* 1193-1221.

A Appendix

Proof of Proposition 1

We use backward induction to solve the game. Consider a firm that observes message $m \in M = \{NR, R\}$ and has an outside candidate x. Let $s_R \equiv E[s|m = R]$ be the firm's updated belief about the expected skill of the focal candidate after observing a recommendation, taking into account Bayes' rule and the adviser's equilibrium strategy. Similarly define the firm's belief after no recommendation, $s_{NR} \equiv E[s|m = NR]$. Hence, s_m is the expected skill of the focal candidate. Hiring the outside candidate yields payoff x to the firm. Therefore, the firm will hire the adviser's candidate if and only if $s_m \geq x$. Since x follows distribution F, after a message m, the adviser's candidate is hired with probability $F(s_m)$.

Now consider the adviser's decision at the beginning of the game. In order to choose a message m, he takes into account the state s, how the message changes the firm's belief and the resulting hiring probability. Moreover, he takes into account how the message can change his reputation.

Let $U(m, s)$ be the adviser's expected utility when he selects message m after observing skill s. If the adviser does not send a recommendation, then there is no reputation change and the firm hires the candidate with probability $F(s_{NR})$. The adviser gains the bias payoff only if the candidate is hired. Therefore,

$$U(NR, s) \equiv bF(s_{NR}) + v(r). \tag{6}$$

If the adviser sends a recommendation, then the firm updates its belief to s_R and hires the candidate with probability $F(s_R)$. If the candidate is hired, the adviser receives the bias payoff b. The change in reputation depends on whether the adviser is lying. If he recommends a high-skilled candidate, then his reputation payoff from a successful hire will be $v(1+g(r))$, therefore

$$\begin{aligned} U(R,s_H) &\equiv F(s_R)\left[b+v(r+g(r))\right]+(1-F(s_R))v(r) \\ &= F(s_R)\left[b+v(r+g(r))-v(r)\right]+v(r) \\ &= F(s_R)\left[b+G(r)\right]+v(r), \end{aligned} \tag{7}$$

where $G(r) \equiv v(r+g(r))-v(r)$ is the reputation payoff gain from a successful hire. If he recommends a low-skilled candidate, then his reputation payoff from an unsuccessful hire will be $v(r-c(r))$, therefore

$$\begin{aligned} U(R,s_L) &\equiv F(s_R)\left[b+v(r-c(r))\right]+(1-F(s_R))v(r) \\ &= F(s_R)\left[b+v(r-c(r))-v(r)\right]+v(r) \\ &= F(s_R)\left[b-L(r)\right]+v(r), \end{aligned} \tag{8}$$

where $L(r)$ is the reputation payoff loss from an unsuccessful hire, defined by (1).

The adviser chooses the message that maximizes his expected utility. We proceed by proving the following intermediate steps.

Step 1: We will prove that, in equilibrium, it must be the case that the firm is strictly more likely to hire the focal candidate after receiving a recommendation, $s_R > s_{NR}$.

By contradiction, suppose that $s_R \leq s_{NR}$. From (6) and (8), this implies that $U(NR,s_L) > U(R,s_L)$ since there is a lying cost $L(r) > 0$ for sending a recommendation when $s = s_L$. Hence, an adviser with a low-skilled candidate would never send a recommendation. From (6) and (7), an adviser with a high-skilled candidate and sufficiently low bias strictly prefers to send a recommendation (to obtain the reputation gain $G(r) > 0$), while advisers with a high bias might not send a recommendation to increase the hiring probability. Therefore, the firm must update its belief to $s_R = s_H$ after observing a recommendation. After not receiving a recommendation, the firm must update its belief s_{NR} to a weighted average of s_H and s_L, which is strictly lower than s_H, a contradiction. Therefore, we must have $s_R > s_{NR}$, which concludes Step 1.

Step 2: We will show that $s_{NR} = s_L$. From Step 1, $s_R > s_{NR}$, which means that the hiring probability is higher when sending a recommendation, $F(s_R) > F(s_{NR})$. Since $G(r) > 0$, it is immediate to verify that an adviser with a high-skilled candidate always strictly prefers to send a recommendation, $U(R,s_H) > U(NR,s_H)$.

Consider now an adviser with a low-skilled candidate. The adviser strictly prefers to lie and send a recommendation if and only if

$$\begin{aligned} U(R, s_L) &> U(NR, s_L) \\ \iff F(s_R)\left[b - L(r)\right] + v(r) &> F(s_{NR})b + v(r) \\ \iff b\left[\frac{F(s_R) - F(s_{NR})}{F(s_R)}\right] &> L(r). \end{aligned} \tag{9}$$

Since $L(r) > 0$, the inequality does not hold if b is sufficiently close to zero. That is, an adviser with a sufficiently low bias will not send a recommendation. Since only advisers with a low-skilled candidate choose to not send a recommendation, the firm must correctly update its belief: it must be the case that $s_{NR} = s_L$.

Step 3: We will now characterize the adviser's cutoff strategy. Using $s_{NR} = s_L$ from Step 2, rewrite (9) as

$$\begin{aligned} U(R, s_L) &> U(NR, s_L) \\ \iff b\left[\frac{F(s_R) - F(s_L)}{F(s_R)}\right] &> L(r). \end{aligned} \tag{10}$$

The adviser will lie if and only if (10) holds. Since $\left[\frac{F(s_R)-F(s_L)}{F(s_R)}\right] > 0$ and $L(r) > 0$, there exists a unique $b^* \in \mathbb{R}_{++}$ such that inequality (10) holds if and only if $b > b^*$. The adviser's behavior can then be characterized by the threshold b^*: the adviser lies and sends a recommendation if his bias is high, $b > b^*$, and he does no send a recommendation if $b \leq b^*$.

Step 4: We now characterize the firm's equilibrium belief s_R, which must be consistent with the adviser's strategy threshold b^*.

With probability p the adviser has a high-skilled candidate and always sends a recommendation. In this case, the firm should increase its belief to s_H. With probability $(1-p)(1-Z(b^*))$ the adviser has a low-skilled candidate, has a high bias $b > b^*$, and sends a recommendation. In this case, the firm should decrease its belief to s_L. Bayes' rule imply that the firm's belief should be the weighted average

$$s_R = \frac{ps_H + (1-p)(1-Z(b^*))s_L}{p + (1-p)(1-Z(b^*))}. \tag{11}$$

Since s_R is a function of the equilibrium threshold b^*, we will abuse notation and use $s_R(b^*)$. Note that $s_R(b^*)$ is a strictly increasing continuous function for $b^* \in (0, 1)$.

Step 5: Finally, we need to guarantee that the beliefs of adviser and firm are consistent with each other.

Combining (10) and (11), we need to find a threshold b^* such that

$$b^* \left[\frac{F(s_R(b^*)) - F(s_L)}{F(s_R(b^*))} \right] = L(r). \tag{12}$$

The LHS of (12) is a strictly increasing, continuous function of $b^* \in (0, 1]$. Hence, the highest possible value for the LHS is

$$\bar{B} \equiv \left[\frac{F(s_H) - F(s_L)}{F(s_H)} \right]. \tag{13}$$

We have two cases. If $L(r) \geq \bar{B}$, then the reputation loss $L(r)$ is weakly greater than the maximum benefit from lying, thus we have a corner solution: $b^* = 1$ and this adviser will not lie. Otherwise, we have an interior solution: there is a unique $b^* \in (0, 1)$ which solves (12). Note that the solution is unique because the LHS is a strictly increasing continuous function of b^*.

It is useful to define a reputation cutoff $\bar{r}$ as follows. If $L(r) < \bar{B}$ for all $r \in \mathbb{R}_{++}$, then define $\bar{r} = \infty$. Otherwise, define $\bar{r}$ to be the unique $r \in \mathbb{R}_{++}$ such that $L(r) = \bar{B}$ (it is unique since $L(r)$ is a strictly increasing continuous function). Hence, if $r \geq \bar{r}$, then we have the corner solution $b^* = 1$. If $r < \bar{r}$, then we have an interior solution and b^* solves (12), concluding the proof. ■

Proof of Proposition 2

The result follows immediately from Proposition 1 and condition (12).

First, if $r \geq \bar{r}$, then the adviser never lies. Further increasing the reputation has no equilibrium effect. Second, if $r < \bar{r}$, then (12) must hold. A marginal increase in r strictly increases the RHS of (12). The LHS of (12) is an increasing function of b^*. Therefore, an increase in r must be followed by an increase in b^*, which implies that the adviser lies less often.

Finally, because the adviser lies less, the firm observes a Blackwell more informative signal[26] from the adviser; hence, the firm's expected payoff goes up. ■

Proof of Proposition 3

From Proposition 2, we know that the adviser lies less if his reputation is higher. The results in Proposition 3 follow from this fact, combined with standard results regarding payoffs under uncertainty.

[26] A signal $S1$ is considered more informative than another signal $S2$ if every Bayesian decision maker prefers $S1$ to $S2$.

The probability of hiring the focal candidate as a function of the firm's beliefs is

$$E\left[F(s_m)\right] \;=\; \Pr(m=R)F(s_R(b^*))+(1-\Pr(m=R))F(s_{NR}). \tag{14}$$

We can interpret F as a utility function and note that the message represents a fair gamble — Bayes' rule implies $\Pr(m=R)s_R(b^*)+(1-\Pr(m=R))s_{NR}=\hat{s}$. The hiring probability $E\left[F(s_m)\right]$ can then be interpreted as an expected utility over this fair gamble.

From Proposition 2, a higher reputation r implies a higher b^*. A higher b^* results in a higher $s_R(b^*)$ and a constant $s_{NR}=s_L$. This represents a riskier fair gamble (a mean-preserving spread). In this case, the expected utility $E\left[F(s_m)\right]$ goes up if F is strictly convex (risk-lover) and goes down if F is strictly concave (risk averse). That is, the hiring probability goes up if the firm has a top-heavy outside market and it goes down if the firm has a bottom-heavy outside market, concluding parts (i) and (ii) of the proposition. Part (iii) directly follows from (i) and (ii), concluding the proof.

■

Proof of Proposition 4

We want to compute the adviser's expected reputation $E[\tilde{r}]$ at the end of the period. With probability p the adviser has a high-skilled candidate and sends a recommendation. Then, the firm hires the candidate with probability $F(s_R(b^*))$ and the reputation becomes $\tilde{r}=r+g(r)$; the firm does not hire the candidate with probability $1-F(s_R(b^*))$ and the reputation stays at $\tilde{r}=r$. With probability $(1-p)$ the adviser has a low-skilled candidate. In this case, with probability $1-Z(b^*)$ he has a high bias and sends a recommendation. The candidate is then hired with probability $F(s_R(b^*))$, resulting in a lower reputation $\tilde{r}=r-c(r)$; with probability $1-F(s_R(b^*))$ the firm does not hire the candidate and the reputation stays at $\tilde{r}=r$. Finally, with probability $Z(b^*)$ the adviser has a low bias and does not send a recommendation, keeping his reputation $\tilde{r}=r$. The adviser's expected reputation at the end of the period is then

$$\begin{aligned}
E\left[\tilde{r}\right] &= p\left\{F(s_R(b^*))(r+g(r))+(1-F(s_R(b^*))r\right\} \\
&\quad +(1-p)\left\{(1-Z(b^*))\left[F(s_R(b^*))(r-c(r))+(1-F(s_R(b^*))r\right]+Z(b^*)r\right\} \\
&= r+F(s_R(b^*))\left\{pg(r)-(1-p)(1-Z(b^*))c(r)\right\} \\
&= r+F(s_R(b^*))\Psi(r),
\end{aligned}$$

where we define $\Psi(r)\equiv\{pg(r)-(1-p)(1-Z(b^*))c(r)\}$. Note that b^* is a continuous, weakly increasing function of r, and $\Psi(r)$ is a continuous function of r. The adviser's expected future reputation $E[\tilde{r}]$ is strictly higher than his current reputation r if $\Psi(r) > 0$, and strictly lower if $\Psi(r) < 0$. Therefore, to prove the proposition, it is sufficient to show that (a) $\Psi(r)$ is strictly positive if r is sufficiently high, (b) $\Psi(r)$ is strictly negative if r is sufficiently low, and (c) $\Psi(r)$ crosses zero at a unique value r_T.

To prove (a), first note that $\bar{B} = \frac{F(s_R(b^*) - F(s_L)}{F(s_R(b^*))} \leq 1$. The proposition assumes that $L(r) > 1$ for some $r \in \mathbb{R}_{++}$, which implies that there exists some $\bar{r} \in \mathbb{R}_{++}$ such that $\bar{B} = L(\bar{r})$. From the proof of Proposition 1, $b^* = 1$ for all $r \geq \bar{r}$. Since $Z(1) = 1$, we have $\Psi(r) = pg(r) + 0 > 0$ for all $r \geq \bar{r}$. Moreover, by continuity and the fact that $c(r) < r$, we also have $\Psi(r) > 0$ for values of r sufficiently close to $\bar{r}$, concluding this step.

To prove (b), recall that b^* and $Z(b^*)$ will be close to zero when r is close to zero. Hence, fixed the function $c(r)$, we need the gain $g(r)$ to be sufficiently small for small values of r, so that $\Psi(r)$ becomes strictly negative. Note that b^* is only a function of $c(r)$, hence reducing $g(r)$ does not change b^*. This concludes this step.

Together steps (a) and (b) imply that the function $\Psi(r)$ crosses zero at least one time from negative to positive. It remains to show that it only crosses once. Let r_T be the highest r such that $\Psi(r) = 0$. From step (a), this point exists and must be below $\bar{r}$, with $\Psi(r) > 0$ for all $r > r_T$. Fixed $c(r)$ and b^*, we need $g(r)$ to be sufficiently low for all $r < r_T$, such that $\Psi(r) < 0$ for all $r < r_T$, concluding the proof.

In the main text, we mention the example $g(r) = \kappa r$ and $c(r) = \delta r$ for some $\kappa > 0$ and $\delta \in (0, 1)$. In this case,

$$\begin{aligned} \Psi(r) &= p\kappa r - (1-p)(1-Z(b^*))\delta r \\ &= r\left\{p\kappa - (1-p)(1-Z(b^*))\delta\right\}. \end{aligned}$$

The term in brackets $\{p\kappa - (1-p)(1-Z(b^*))\delta\}$ is an increasing function of r, since b^* is an increasing function of r. When r goes to zero, the lying probability $(1 - Z(b^*))$ goes to one, and the term in brackets goes to $p\kappa - (1-p)\delta$. Hence, we need to assume that $\kappa < \frac{(1-p)\delta}{p}$, so that $\Psi(r) < 0$ for low values of r. Given this κ, when r is sufficiently high, the lying probability $(1 - Z(b^*))$ goes to zero, and the term in brackets always becomes strictly positive. Finally, since the term in brackets is strictly increasing for $r < \bar{r}$, it crosses zero exactly one time. ■

Proof of Proposition 5

Proof of Part (i): Consider outside markets F and $\hat{F}$, with $F(x) = \hat{F}(x)$ for all $x \in (0, \hat{s})$, $F(x) > \hat{F}(x)$ for all $x \in (\hat{s}, 1)$, and interior equilibrium $b^* \in (0, 1)$. An interior equilibrium implies that (12) holds. Note that $s_R(b^*) > \hat{s} > s_L$. Holding b^* constant, the change in F implies $F(s_R(b^*)) > \hat{F}(s_R(b^*))$ and $F(s_L) = \hat{F}(s_L)$. Therefore, (12) and the change in F imply that, if we hold constant b^*, we have the strict inequality[27]

$$b^* \left[\frac{\hat{F}(s_R(b^*)) - \hat{F}(s_L)}{\hat{F}(s_R(b^*))} \right] \;<\; L(r). \tag{15}$$

Since the LHS is an increasing function of b^*, we need to increase b^* to achieve the new equilibrium. A higher b^* means that the adviser lies less. The firm benefits twice: it directly benefits from a better pool of outside candidates, and indirectly because the adviser lies less, concluding this part of the proof.

Proof of Part (ii): Consider outside markets F and $\hat{F}$, with $F(x) = \hat{F}(x)$ for all $x \in (\hat{s}, 1)$, $F(x) > \hat{F}(x)$ for all $x \in (0, \hat{s})$, and interior equilibrium $b^* \in (0, 1)$. An interior equilibrium implies that (12) holds. Note that $s_R(b^*) > \hat{s} > s_L$. Holding b^* constant, the change in F implies $F(s_R(b^*)) = \hat{F}(s_R(b^*))$ and $F(s_L) > \hat{F}(s_L)$. Therefore, (12) and the change in F imply that, if we hold constant b^*, we have the strict inequality[28]

$$b^* \left[\frac{\hat{F}(s_R(b^*)) - \hat{F}(s_L)}{\hat{F}(s_R(b^*))} \right] \;>\; L(r). \tag{16}$$

Since the LHS is an increasing function of b^*, we need to decrease b^* to achieve the new equilibrium. A lower b^* means that the adviser lies more.

The firm faces a tradeoff: it directly benefits from better outside candidates, but indirectly loses from the adviser lying more. Which effect dominates depends on the parameters of the model. We next construct an example in which the firm's payoff goes down when we improve the distribution of bottom-candidates. The key is to change F as follows: move probability mass from just below s_L to just above it. This decreases $F(s_L)$ significantly, leading the adviser to lie significantly more, with just a small effect on the expected skill of the alternative candidate.

Example: Suppose $r < \bar{r}$, which implies that (12) holds, $b^* \in (0, 1)$, and $0 < s_L < \hat{s} <$

[27] The LHS of (12) can be rewritten as $b^* \left[1 - \frac{F(s_L)}{F(s_R(b^*))}\right]$. Holding b^* constant, the change in F decreases the denominator from $F(s_R(b^*))$ to $\hat{F}(s_R(b^*))$, while holding constant the numerator $F(s_L)$. Consequently, the LHS decreases.

[28] The LHS of (12) can be rewritten as $b^* \left[1 - \frac{F(s_L)}{F(s_R(b^*))}\right]$. Holding b^* constant, the change in F decreases the numerator from $F(s_L)$ to $\hat{F}(s_L)$, while holding constant the denominator $F(s_R(b^*))$. Consequently, the LHS increases.

$s_R(b^*) < 1$. The firm's expected payoff is

$$U^{DM} = \Pr(m{=}R)\Gamma_R + (1 - \Pr(m{=}R))\Gamma_{NR}, \tag{17}$$

where the firm's expected payoff conditional on receiving a recommendation is

$$\Gamma_R = s_R(b^*)F(s_R(b^*)) + \int_{s_R(b^*)}^{1} xf(x)dx, \tag{18}$$

and the firm's expected payoff conditional on not receiving a recommendation is

$$\Gamma_{NR} = s_L F(s_L) + \int_{s_L}^{1} xf(x)dx. \tag{19}$$

Construct an outside market $\hat{F}$ with better bottom-candidates as follows: move a small probability mass ϵ from $(0, s_L)$ to $(s_L, \hat{s})$. That is, keep the same distribution on the top, $F(x) = \hat{F}(x)$ for all $x \in [\hat{s}, 1]$, decrease the pdf below s_L, $f(x) > \hat{f}(x)$ for all $x \in (0, s_L)$, and increase the pdf above s_L, $f(x) < \hat{f}(x)$ for all $x \in (s_L, \hat{s})$, where $\epsilon = F(s_L) - \hat{F}(s_L)$ is the small probability mass that was moved from the range $(0, s_L)$ to $(s_L, \hat{s})$.

We will now determine how a marginal change in ϵ affects the firm's payoff. Marginally increasing ϵ from $\epsilon = 0$ is an improvement of bottom-candidates, which leads the adviser to lie more. This does not change s_L but it decreases $s_R(b^*)$ and increases $\Pr(m{=}R)$. Moreover, it only changes F and f below $\hat{s}$. Take the derivative

$$\begin{aligned}\frac{\partial U^{DM}}{\partial \epsilon} &= \frac{\partial \Pr(m{=}R)}{\partial \epsilon}\Gamma_R + \Pr(m{=}R)\frac{\partial \Gamma_R}{\partial \epsilon} - \frac{\partial \Pr(m{=}R)}{\partial \epsilon}\Gamma_{NR} + (1 - \Pr(m{=}R))\frac{\partial \Gamma_{NR}}{\partial \epsilon} \\ &= \frac{\partial \Pr(m{=}R)}{\partial \epsilon}[\Gamma_R - \Gamma_{NR}] + \Pr(m{=}R)\frac{\partial \Gamma_R}{\partial \epsilon} + (1 - \Pr(m{=}R))\frac{\partial \Gamma_{NR}}{\partial \epsilon}. \end{aligned} \tag{20}$$

Since $\Pr(m{=}R) = \frac{\hat{s} - s_L}{s_R(b^*) - s_L}$, we have

$$\frac{\partial \Pr(m{=}R)}{\partial \epsilon} = -\frac{\hat{s} - s_L}{(s_R(b^*) - s_L)^2}\frac{\partial s_R(b^*)}{\partial \epsilon} = -\frac{\Pr(m{=}R)}{(s_R(b^*) - s_L)}\frac{\partial s_R(b^*)}{\partial \epsilon}. \tag{21}$$

Moreover

$$\begin{aligned}\frac{\partial \Gamma_R}{\partial \epsilon} &= \frac{\partial s_R(b^*)}{\partial \epsilon}F(s_R(b^*)) + s_R(b^*)f(s_R(b^*))\frac{\partial s_R(b^*)}{\partial \epsilon} \\ &\quad - s_R(b^*)f(s_R(b^*))\frac{\partial s_R(b^*)}{\partial \epsilon} = \frac{\partial s_R(b^*)}{\partial \epsilon}F(s_R(b^*))\end{aligned} \tag{22}$$

and

$$\Gamma_R - \Gamma_{NR} = s_R(b^*)F(s_R(b^*)) - s_L F(s_L) - \int_{s_L}^{s_R(b^*)} x f(x)dx. \quad (23)$$

Substitute (21), (22) and (23) into (20) and simplify to obtain

$$\begin{aligned}
\frac{\partial U^{DM}}{\partial \epsilon} &= \frac{\Pr(m{=}R)}{(s_R(b^*) - s_L)} \frac{\partial s_R(b^*)}{\partial \epsilon} \Bigg\{ - s_R(b^*)F(s_R(b^*)) + s_L F(s_L) + \int_{s_L}^{s_R(b^*)} x f(x)dx \\
&\quad + (s_R(b^*) - s_L)F(s_R(b^*)) \Bigg\} + (1 - \Pr(m{=}R)) \frac{\partial \Gamma_{NR}}{\partial \epsilon} \\
&= \frac{\Pr(m{=}R)}{(s_R(b^*) - s_L)} \frac{\partial s_R(b^*)}{\partial \epsilon} \left\{ \int_{s_L}^{s_R(b^*)} x f(x)dx - s_L(F(s_R(b^*)) - F(s_L)) \right\} \\
&\quad + (1 - \Pr(m{=}R)) \frac{\partial \Gamma_{NR}}{\partial \epsilon} \\
&= \frac{\Pr(m{=}R)}{(s_R(b^*) - s_L)} \frac{\partial s_R(b^*)}{\partial \epsilon} \left\{ \int_{s_L}^{s_R(b^*)} (x - s_L) f(x)dx \right\} + (1 - \Pr(m{=}R)) \frac{\partial \Gamma_{NR}}{\partial \epsilon} \quad (24)
\end{aligned}$$

The term $\frac{\partial s_R(b^*)}{\partial \epsilon}$ is strictly negative, since the improvement in the distribution of bottom-candidate causes the adviser to lie more. The expression

$$\frac{\Pr(m{=}R)}{(s_R(b^*) - s_L)} \frac{\partial s_R(b^*)}{\partial \epsilon} \left\{ \int_{s_L}^{s_R(b^*)} (x - s_L) f(x)dx \right\}$$

captures the firm's loss from the adviser lying more. This loss depends on how much probability mass ϵ we move from $(0, s_L)$ to $(s_L, \hat{s})$, but it does not depend on the particular changes in f. That is, fixed ϵ, we are free to shift the pdf within each range. To see this, recall the equilibrium condition (12), $b^* \left[\frac{F(s_R(b^*)) - F(s_L)}{F(s_R(b^*))} \right] = L(r)$. Since our change does not affect F above the prior and s_L is fixed, the new equilibrium condition $\hat{b}^*$ is

$$\begin{aligned}
\hat{b}^* \left[\frac{F(s_R(\hat{b}^*)) - \hat{F}(s_L)}{F(s_R(\hat{b}^*))} \right] &= L(r) \\
\Rightarrow \hat{b}^* \left[\frac{F(s_R(\hat{b}^*)) - F(s_L) + \epsilon}{F(s_R(\hat{b}^*))} \right] &= L(r),
\end{aligned}$$

where we used the fact that $\hat{F}(s_L) = F(s_L) - \epsilon$.

Since the first term in (24) is strictly negative, we need to guarantee that the second term $(1 - \Pr(m{=}R)) \frac{\partial \Gamma_{NR}}{\partial \epsilon}$ is sufficiently close to zero, so that (24) is negative. The derivate $\frac{\partial \Gamma_{NR}}{\partial \epsilon}$

captures the increase in the firm's payoff conditional on not receiving a recommendation. We can make $\frac{\partial \Gamma_{NR}}{\partial \epsilon}$ arbitrarily small by concentrating the increase ϵ close to s_L. To see this, suppose all increase is placed at s_L: that is, we move probability mass ϵ from $(0, s_L)$ and place all of it at s_L. Then $\frac{\partial \Gamma_{NR}}{\partial \epsilon} = 0$. By continuity, if we spread the probability mass ϵ in the interval $(s_L, \hat{s})$ but place a sufficiently large share of it arbitrarily close to s_L, then we can guarantee that $\frac{\partial \Gamma_{NR}}{\partial \epsilon}$ is close to zero. Hence, fixed ϵ, the first term in (24) is strictly negative and the second term is arbitrarily close to zero if the change in F is concentrated close to s_L. This is then an example in which the firm's payoff goes down after it has access to better bottom-candidates, concluding the proof. ■

Proof of Proposition 6

An interior equilibrium implies that (12) holds. Note that $1 > s_R(b^*) > \hat{s} > s_L > 0$. Holding b^* constant, the rotation of F around $\hat{s}$ implies $F(s_R(b^*)) > \hat{F}(s_R(b^*))$ and $F(s_L) < \hat{F}(s_L)$. Therefore, (12) and the change in F imply that, if we hold b^* constant, we have the strict inequality[29]

$$b^* \left[\frac{\hat{F}(s_R(b^*)) - \hat{F}(s_L)}{\hat{F}(s_R(b^*))} \right] \quad < \quad L(r). \tag{25}$$

Since the LHS is an increasing function of b^*, we need to increase b^* to achieve the new equilibrium. A higher b^* means that the adviser lies less, concluding the proof. ■

[29]The LHS of (12) can be rewritten as $b^* \left[1 - \frac{F(s_L)}{F(s_R(b^*))}\right]$. Holding b^* constant, the change in F strictlty increases the numerator from $F(s_L)$ to $\hat{F}(s_L)$, and strictly decreases the denominator from $F(s_R(b^*))$ to $\hat{F}(s_R(b^*))$. Consequently, the LHS decreases.