

“The Role of Referrals as Matchmakers: Theory”

Lalith Munasinghe, Kathryn Gautier

Recruiters and hiring managers are tasked with one of the most important people management roles in a company: finding and hiring the best talent. Their goal may be quality (attracting and selecting productive and durable employees), quantity (keeping up with demand to fill seats), or both. The challenge is of course identifying the “best talent” with limited information available at the time of recruitment and hire.

Companies try to resolve this information problem through interviews, screening, and testing, among other strategies. Referral networks, both formal and informal, play a major role in the resolution of uncertainty in forming new employment relationships.

Evidence abounds—in industry reports, business publications, and rigorous labor economic research—that referred applicants and employees are of higher quality. In particular, referred applicants are more likely to accept an offer, reducing time to hire; more likely to stay at the job longer, reducing turnover; and more likely to perform better, increasing overall productivity. A better understanding of this talent source therefore presents an opportunity to improve the recruitment process and quality of employees. A recent research paper provides a simple but powerful theory, including supporting evidence, to explain the importance of referrals: information¹.

In this white paper, we introduce a basic framework to understand the role of information in the referral process, showing that the information introduced by referrals leads to fewer costly hiring errors (“quick quits”), and the identification of high quality employees who would otherwise go overlooked in the traditional recruiting processes. These positive outcomes can be achieved only by tracking important referral information, and by leveraging the right types of referrals. This paper is the first in a two-part series covering referrals; the second explores the empirical implications of the theory outlined in this paper, using data from a large customer services company.

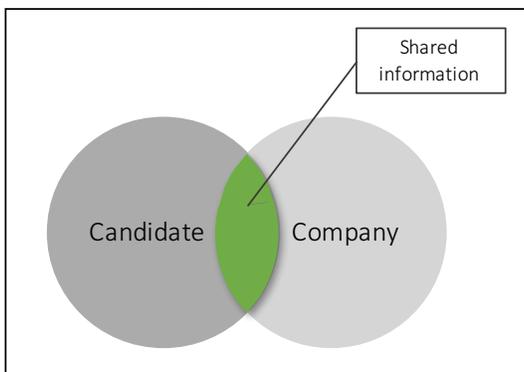


Figure 1.1. The candidate and company share some set of information about each other, from the application, testing, interviews, public resources, etc. They use this shared information, shown in green, to form assessments about potential match quality.

Summary

- Hiring decisions are based on expectations about employment match quality
- Errors—hiring the wrong person, or missing out on the right person—are due in part to lack of information about how good a match will be
- Referrals introduce information that enable better hiring decisions
 - Avoid “quick quits”
 - Identify and take chances on “diamonds in the rough”
- Incentivize referrers to share accurate, complete information, and thus generate a pipeline of more durable and productive talent

To clarify the role that information plays in recruitment, consider the typical hiring process in the absence of a referral. A candidate may have some information about a job, from the published description of the open position, materials about the company available online, etc. Similarly, the company will have some information about the candidate,

¹ Barr, T., Bojilov, R., and Munasinghe, L. “Referrals and Search Efficiency: Who Learns What and When?” Forthcoming in Journal of Labor Economics. 2017.



from her application, resume, references, etc. (Figure 1.1). Throughout the application and interview process, both the candidate and the firm aim to discover the same thing: *would the job be a good fit for the candidate; and would the candidate be a good fit for the job?* In not so many words: match quality. The idea of being a *good match* can encompass many different positive outcomes for the firm and the employee, from increased productivity on the job to greater employee satisfaction and happiness, that typically result in a more durable employment relationship. Note that a greater overlap—green area of Fig. 1.1—does not mean that a candidate and company are better matched, simply that they are in a better position to *identify* if they are (or are not) a good match.

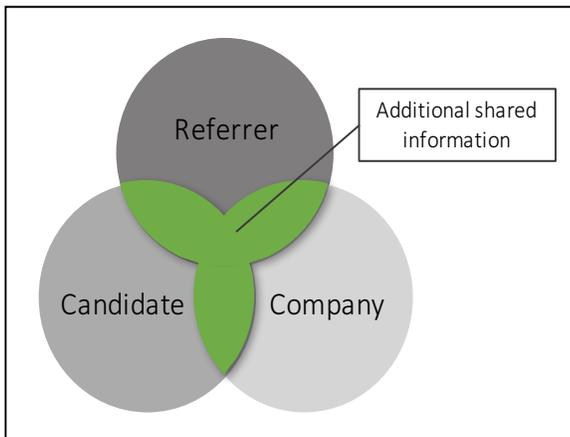


Figure 1.2. The referrer has additional match-relevant information, through his experience working for the company, and his relationship with the candidate, thus increasing the overlap of available information.

Even after this initial screening process, there still remains, of course, some uncertainty about the quality of the employment match. However, both parties must make decisions in the face of this uncertainty, possibly resulting in errors: There may be both Type I errors—badly-matched candidates entering the job—and Type II errors—well-matched candidates failing to enter the job—due at least in part to the lack of information at this early decision point.

Reducing the errors in these decisions of course would be in the interest of both the company and the candidate, however obtaining additional match-relevant information can be costly. Referrals can be viewed as a cost-effective strategy to elicit more of this information.

the potential employment match, and the candidate and company exchange information through an application, interview, etc. But in this case, there is also a third party—the referrer—who has additional information both about the company where he works, and the candidate with whom he is acquainted (Figure 1.2). If he shares that information with the respective parties, both will be in a better position to resolve some of the information uncertainty and thus make more informed decisions². The result of this additional exchange of information is fewer hiring errors, both Type I and Type II, implying more productive and durable hires.

Consider the same screening process, in the presence of a referral. As before, the goal is to determine the quality of

Within this basic theoretical framework of information exchange, we can now explore the implications that referrals will have for recruiting and hiring. Let's first formalize hiring and employment in terms of the more- or less-informed decisions discussed earlier. We break down the hiring and employment process into five stages, resulting in five groups of people, each a subset of the previous: potential candidates for the job, those who apply, those offered a position, those

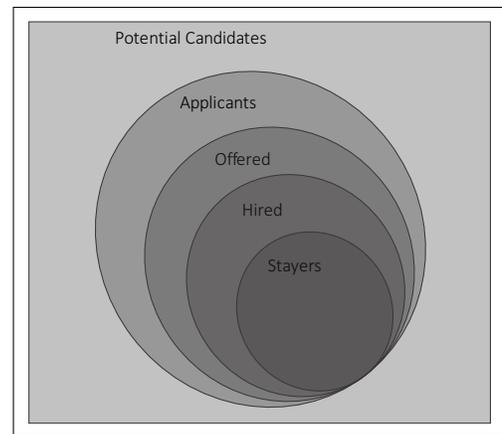


Figure 2. Five stages of the hiring and recruiting process.

² We assume here that the referrer freely shares the relevant information with both the firm and the candidate, but the revelation of information is a nontrivial issue, and in fact incentives must be present for the referrer to share both accurate and complete information. It may also be the case that the referrer shares information only with the firm, and indeed this may be in the firm's best interest. These points introduce a complexity outside the scope of this paper, but are more carefully explored in the referenced paper by Barr, Bojilov, and Munasinghe and in the broader body of literature on this topic.



who accept and are hired, and finally those who stay at the job for an extended period, at which point we assume the information uncertainty is resolved (Figure 2).³

We then think of the movement of a person from one stage to the next as a *selection decision*: for a given match, one or both parties decide to select in or out at each stage of this job matching process. Selecting in to the next stage is driven by expected match quality: given what each party knows about the other at that time, how well-matched do they think they are? This expected match quality must be large enough to outweigh the expected costs associated with each stage (e.g. the time it takes to fill out an application, the cost of onboarding a new employee, or the investment in training).⁴

At each stage, selection decisions are based on the information available at that time. So, if there is match-relevant information, we would expect the average match quality to improve after each decision, because in general poorly-matched candidates will tend to select out, and well-matched candidates will tend to select in (though there are of course errors). The greater the quality and quantity of information, the more we expect average match quality to improve at each stage, due to prevention of errors in selection decisions.

Therefore, due to the additional information referrals bring to bear in the hiring process, the pre-hire selection decisions made in the presence of a referral are better. In other words, there are fewer selection errors in the referred sample, and so there is greater improvement between each group. Figure 3 shows qualitatively this process separated for referred and non-referred candidates.

Average match quality is measured on the x-axis, for each of the samples of people. The hypothesized relative size of each pool is measured on the y-axis. For expositional purposes, we assume that the average quality of potential candidates, and the number of applicants is the same across referred and non-referred. The distance between each subsequent group on the x-axis represents the improvement in average match quality after each stage of selection decisions; the drops on the y-axis represent the fraction of the group who select out of the next stage of the employment process, due to their assessments of poor match-quality.

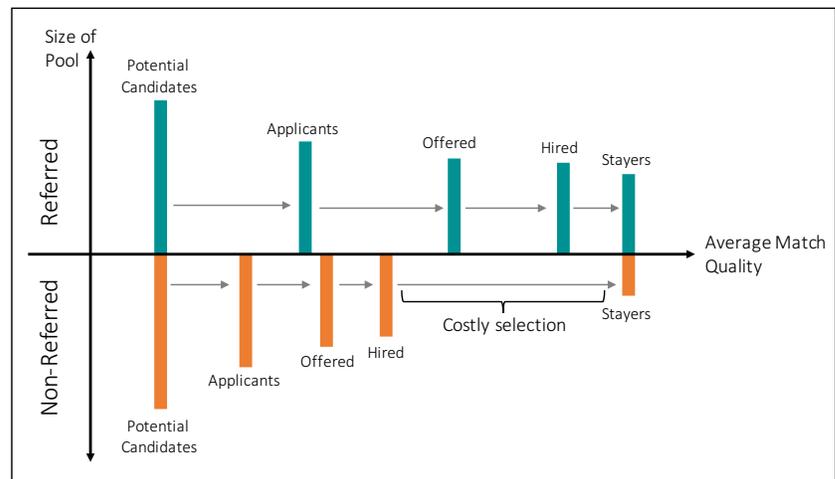


Figure 3. A sketch of average match quality tracked throughout the selection processes, with and without the presence of a referral. The information revealed through the referral process speeds up the sorting, resulting in lower costly on-the-job turnover, and fewer overlooked good matches.

Note that once the candidate is hired, the information implicated in the decision of the candidate and/or the firm to maintain the employment relationship is primarily through first-hand experience. Thus, once the non-referred hires begin working, they directly obtain information related to match quality, select in or out

³ By potential candidates we don't mean the population in general, but rather the set of people who have the requisite skills to perform the job in question, and who are in the relevant and feasible (e.g. local) labor market.

⁴ Clearly, the greatest costs are incurred by both parties at the post-hire stage, when there is investment in training. Thus, we often think of early on-the-job turnover as a major pain point, since most of these costs have already been borne by both parties, but few benefits have been reaped by either.



of employment accordingly, and ultimately their average quality is thus presumed identical to that of the referred stayers: the only difference is *when* they obtained the information.

This differential timing of information revelation is of practical importance for two reasons.

The first is that post-hire, on-the-job selection is highly costly both for the firm and the employee; it is much more favorable to make the correct decision to select out of employment before hire. Referrals allow for just that, reducing on-the-job turnover by introducing match-relevant information to the application process.

The second is that referrals identify good matches that would otherwise have been missed, for instance because the candidate never would have applied, or because the company ordinarily would have rejected based on the traditionally available information. In technical terms, the referrer pool generates a greater quantity of good matches (stayers) from the same number of applicants. This is because of the irreversibility of Type II errors (good matches incorrectly selecting out of employment). While all Type I errors (bad matches selecting in) are eventually corrected through on-the-job turnover as all information is revealed, Type II errors that occurred pre-hire cannot be corrected. Since Type II errors occur with greater frequency in the non-referred pool, there will be fewer good matches remaining even after the uncertainty about match quality is resolved⁵. The identification of these so-called “diamonds in the rough” are perhaps the greatest unforeseen benefit of referrals from a qualitative perspective.

These results hold under our assumption that the referred candidates pool is *no better* than the non-referred pool; outcomes here are due simply to the additional information revealed through the referral process.

If the value of referrals is really in the information they bring to bear, then the question of how well-informed is a given referrer becomes crucial. A referrer may be well-informed about the company—has worked there a long time, has worked in several different roles—or he may be less informed about the

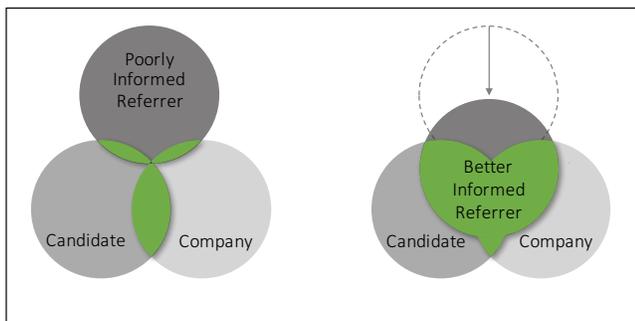


Figure 4. With better informed referrers, there is even greater overlap in shared information, enabling more accurate decision making.

company—recently hired or a narrow work history. Similarly, a referrer may be more or less informed about the candidate—they could have worked together in the past, or perhaps they simply share a mutual friend. This differential information content across referrals (Figure 4) will have consequences for the selection process and employment outcomes.⁶ In short, with better informed decision makers, outcomes will be better (due to the selection process outlined earlier).

Hence, the understanding of referrals and the nature of the referral relationships becomes critical. In

making people decisions, whether aided by data and analytics or not, knowing from which *type* of referral a candidate comes is key in predicting his or her future success. It’s not enough to simply encourage referrals in general, rather it’s necessary to properly incent the *right* referrers. These key pieces of additional

⁵ Note that we can say nothing definitive about the relative sizes of the offered and hired pools because the Type I and Type II errors act in opposite directions, and it is not clear which will dominate the other.

⁶ In addition to differential information content, different employees may also have differential network qualities, e.g. a top performer may know people with similar qualities. This will have further implications for employment outcomes.

information have historically been ignored in both academic literature and industry practices, but they should be collected and systematically leveraged to create better policies, insights, and decisions.

The main ideas in this paper are not specific to referrals. Rather, the role of information uncertainty in job matching is a broad subject of much interest in labor economics, and the theory surrounding these ideas (called search theory, or matching theory) is applicable throughout the recruiting, hiring, and employment process⁷. Just as quality referrals can introduce additional relevant information and facilitate better decision making, there are several further sources of information which can improve employment outcomes, for example, quality interviewers, intelligently designed applications, and systematic tracking of historical outcomes.

Understanding how exactly referrals and these other sources of information generate better quality employees is crucial to generating intelligent recruiting policies, smart data collection, and coherent incentive schemes.

⁷ Munasinghe, Lalith, 2006. Expectations matter: Job prospects and turnover dynamics. *Labour Economics*, 5(13), pp.589-609.

