

The Heterogeneous Impact of Referrals on Labor Market Outcomes

Benjamin Lester

Federal Reserve Bank of Philadelphia Research Department

David A. Rivers

University of Western Ontario

Giorgio Topa

Federal Reserve Bank of New York and IZA

WP 21-34

PUBLISHED

October 2021

REVISED

August 2024

ISSN: 1962-5361

Disclaimer: This Philadelphia Fed working paper represents preliminary research that is being circulated for discussion purposes. The views expressed in these papers are solely those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of Philadelphia or the Federal Reserve System. Any errors or omissions are the responsibility of the authors. Philadelphia Fed working papers are free to download at: <https://philadelphiafed.org/research-and-data/publications/working-papers>.

DOI: <https://doi.org/10.21799/frbp.wp.2021.34>

The Heterogeneous Impact of Referrals on Labor Market Outcomes^{*†}

Benjamin Lester

Federal Reserve Bank of Philadelphia

David A. Rivers

University of Western Ontario

Giorgio Topa

Federal Reserve Bank of New York and IZA

February 8, 2024

Abstract

We study the impact of referrals on labor market outcomes. First, we document a new set of facts exploiting data that allow us to distinguish directly between different types of referrals—those from family and friends and those from business contacts—and different types of jobs, as measured by the skill requirements of the occupation. We then develop a structural framework to interpret these facts and quantify the effects of social and business networks on employment, earnings, output, and inequality. Referrals from family and friends generate good jobs for all workers but are relied upon by those who struggle to generate offers through other channels. Referrals from business contacts are used predominantly by more productive workers who receive offers through other channels relatively frequently. An important implication is that referrals from business contacts exacerbate earnings inequality, while referrals from family and friends actually reduce inequality.

Keywords: Labor Markets, Referrals, Networks, Search Theory, Asymmetric Information

JEL Classification: E42, E43, E44, E52, E58

^{*}We thank Assa Cohen, Linda Wang, and Junyuan Zou for exceptional research assistance. For helpful comments, we thank Audra Bowlus, Jason Faberman, Shigeru Fujita, Manolis Galenianos, Greg Kaplan, Rasmus Lentz, Igor Livshits, Lance Lochner, Ryan Michaels, Simon Mongey, Giuseppe Moscarini, Ian Schmutte, and David Wiczer, along with conference and seminar participants at the NBER Summer Institute, the Annual Meeting of the Society for Economic Dynamics, the DC Search and Matching Workshop, the Workshop on Market Frictions and Macroeconomics, the Barcelona School of Economics Summer Forum, Boston College, Federal Reserve Bank of St. Louis, Johns Hopkins University, University of Toronto, Vanderbilt University, Wharton, University of California–Irvine, University of Georgia, University of Western Ontario, and Waterloo University. All errors are our own.

[†]**Disclaimer:** The views expressed in this paper are solely those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of Philadelphia, the Federal Reserve Bank of New York, or the Federal Reserve System. Any errors or omissions are the responsibility of the authors.

1 Introduction

According to surveys of workers and firms, a *referral* is used somewhere in the hiring process for approximately half of all jobs.¹ Given its prevalence, this feature of job search is presumably an important factor in determining how quickly workers and firms form matches, the quality of these matches, and the subsequent implications for wages and turnover. Moreover, as access to (and reliance on) referral networks tends to be heterogeneous across segments of the labor force, the use of referrals has the potential to generate or ameliorate economic inequality.²

Despite the prevalence of referrals and their potentially important implications for labor market outcomes, it remains unclear exactly what (if anything) referrals actually do. While the theoretical literature has proposed a variety of roles that referrals could play in the match formation process, understanding and quantifying the effects of referrals has proven difficult, for at least two reasons. First, there are few representative datasets that contain direct information about whether a referral was used in the hiring process. Second, even when referrals are observed in the data, the set of workers who use referrals is likely to be a selected sample. In light of these challenges, it is perhaps not surprising that the existing empirical literature has found mixed evidence regarding even basic facts about referrals, such as the types of workers that use referrals most frequently or the effects of a referral on a worker’s starting wage.³

In this paper, we first document a number of new facts about the types of workers and jobs that use a referral to form a new match, and how the use of a referral correlates with important labor market outcomes, such as wages and turnover. To do so, we exploit a relatively new survey that contains direct information about whether workers used a referral and, crucially, enables us to distinguish between *different types of referrals*—namely, those from family and friends and those from business contacts—and *different types of jobs*, as measured by the skill requirements of the occupation.

We find that referrals from business contacts are used relatively more frequently to form

¹Topa (2011) provides an extensive review of usage rates across surveys of both workers and firms. Most surveys of job seekers find between 50 and 60 percent of workers report using a referral to find employment (Corcoran et al., 1980; Lin et al., 1981; Bridges and Villemez, 1986; Granovetter, 1995), though others find even higher usage rates (Holzer, 1987b; Elliott, 1999). Similar rates have also been documented in other countries (Gregg and Wadsworth, 1996; Alon and Stier, 2019; Wahba and Zenou, 2005). Surveys of firms also indicate widespread use of referrals or word-of-mouth techniques, though results vary from just under 40 percent of hires using a referral (Holzer, 1987a; Marsden, 2017) to significantly more than 50 percent (Neckerman and Kirschenman, 1991; Miller and Rosenbaum, 1997).

²See, e.g., the influential work of Calvó-Armengol and Jackson (2004). Ioannides and Loury (2004) provide an extensive overview of the relationship between networks and inequality.

³For example, examining the relationship between the use of a referral and match quality (typically measured using wages), Corcoran et al. (1980), Datcher (1983), Simon and Warner (1992), Marmaros and Sacerdote (2002), Kugler (2003), Bayer et al. (2008), and Dustmann et al. (2016) find a positive relationship; Pistaferri (1999), Mouw (2003), and Bentolila et al. (2010) find a negative relationship; and Marsden and Gorman (2001), Loury (2006), and Pellizzari (2010) report mixed results.

matches in high-skill occupations, and that they are associated with higher starting wages and more turnover (or shorter tenure). In contrast, referrals from family and friends are used more frequently in low-skill occupations, and they are associated with lower starting wages and longer tenure. A key factor in explaining these facts appears to be unobserved heterogeneity in the intensity with which workers make contact with firms, both on- and off-the-job: workers who used a referral from a business contact to find their current job receive offers at a relatively high rate, whereas workers who used a referral from family and friends struggle to generate offers.

To interpret these empirical findings and uncover the underlying economic mechanisms at play, we develop a structural model with multiple job search channels, on-the-job search, and unobserved heterogeneity. Calibrating the model to key moments from our data reveals qualitative insights into the distinct roles played by different types of referrals during the match formation process, along with quantitative estimates of the contribution of these two types of referrals to employment rates, wages, turnover, earnings inequality, and output.

Qualitatively, we find that different theories are required to understand these two types of referrals. In particular, referrals from business contacts primarily screen workers based on their *ex ante* expected productivity (or “type”), as in theories in which referrals ameliorate asymmetric information about workers’ productivity.⁴ In contrast, referrals from family and friends generate good matches for all worker types, as in theories in which referrals improve match quality *ex post* by resolving symmetric uncertainty about match productivity or by easing inefficiencies associated with moral hazard.

Quantitatively, we find that referrals from family and friends provide a key source of earnings for workers at the lower end of the income distribution, particularly in low skill markets, and thus tend to *reduce earnings inequality*. In contrast, referrals from business contacts primarily help high productivity workers with high incomes, particularly in high skill occupations, and thus *exacerbate earnings inequality*. Hence, while both types of referrals are an important source of new and better matches, the oft-discussed trade-off between output and inequality applies only to referrals from business contacts in our calibrated model.

We now describe in greater detail the data we use and the empirical facts we uncover; the model we construct to interpret and analyze these empirical findings; and the main qualitative and quantitative insights that emerge from calibrating this model and using it to perform counterfactual exercises. Then, we explain how these results fit into the existing literature.

Facts. After describing the data in Section 2, we document a number of new facts about the types of workers and occupations that tend to use referrals and the characteristics of matches formed through referrals (relative to those formed through other channels) in Section 3. Our data, which come from the Job Search supplement to the Survey of Consumer Expectations,

⁴We discuss existing theories about the role of referrals in more detail in Section 1.1.

has a number of unique features that make it well-suited to studying these issues: the survey draws from a wide range of demographic groups, industries, and occupations; it contains a rich set of information describing the job characteristics of currently employed workers; and, most importantly, it paints a detailed picture of the job search process that generated the current job (as well as other offers), including *direct information about whether a referral was used in the hiring process*.

In fact, not only does the survey collect information about whether a referral was used in the hiring process, but it also asks the worker to describe *who provided the referral*. We group the responses into two classifications—those from “family and friends” and those from “business contacts”—which allows us to distinguish between referrals drawn from a worker’s social and professional networks, respectively. We find that these two types of referrals are completely different animals, and distinguishing between them is crucial for understanding the role of referrals in the labor market.

To start, the two types of referrals appear to be used more frequently in different markets: referrals from family and friends are used most frequently to form matches in low skill occupations, whereas referrals from business contacts are used more intensely in high skill (relative to low skill) markets. Moreover, the correlations between the use of these two types of referrals and various labor market outcomes are opposites: workers who used a referral from a business contact tend to earn higher starting wages than non-referred workers but experience shorter job tenures; whereas workers who used a referral from a friend or relative to get their current job tend to have lower starting wages than non-referred workers but experience less job turnover. Digging deeper, we find that the most likely explanation for these patterns derives from selection on unobserved heterogeneity. In particular, we document that the workers who got their current job using a business referral tend to meet firms at a relatively high rate—while unemployed and employed—whereas workers who got their current job through their network of family and friends generate contacts with firms at a relatively low rate.

Hence, the first part of the paper establishes that clear empirical relationships between referrals and labor market outcomes emerge once we distinguish along two relatively unexplored dimensions of the data: the source of the referral, and the type of job or occupation. These relationships are not only suggestive about the different roles that the two types of referrals play in the hiring process, but they also help explain why previous studies—which could not make the same distinctions in the data—found mixed or conflicting results.

Model. In Section 4, we introduce a structural model to interpret and analyze the new empirical findings discussed above. We argue that there are (at least) two important reasons to do so. First, while any number of stories could qualitatively fit a subset of the empirical patterns we document, an equilibrium model allows us to identify and test various mechanisms across mul-

tuple, inter-related dimensions. Hence, studying the empirical findings discussed above through the lens of a model reveals new insights into the (distinct) roles played by referrals drawn from workers’ business and social networks. Second, calibrating the model to our micro-data via indirect inference allows us to identify certain parameters and relationships that we simply cannot observe in existing data. Using these estimates of the model’s structural parameters reveals new quantitative insights into the contribution of workers’ social and business networks to earnings, output, and inequality.

The key ingredients of the model are motivated by the patterns we observe in the data. First, given the clear correlations we find between various labor market outcomes and the channel through which a worker found a job, we allow for different technologies for initiating contact between workers and firms. In particular, we let contacts arrive through formal search methods, through referrals from business contacts, or through referrals from family and friends. Moreover, we allow the quality of the match (i.e., the match-specific productivity of a worker-firm pair) to depend on the channel through which the contact was initiated. Second, since some workers appear to generate offers through different channels at different rates, even after controlling for a variety of observable characteristics, we allow for worker heterogeneity along some intrinsic type or “ability.” We allow a worker’s ability to affect both the (exogenous) rate at which the worker meets firms through different channels, along with the match-specific productivity the worker draws conditional on meeting a firm through a specific channel. Finally, since a worker’s ability (or proclivity) to contact firms and generate offers does not evaporate after forming a match, but rather seems to be an important feature of understanding heterogeneity in wages and turnover, we assume that workers search when both unemployed and employed.

Our model can be seen as a natural extension of the workhorse models of on-the-job search with unobserved worker heterogeneity (Postel-Vinay and Robin, 2002; Cahuc et al., 2006), extended to allow for multiple job search channels. The model is intentionally constructed to be rich enough to confront the facts we uncover in Section 3, yet equally as tractable as its predecessors. We exploit this tractability to derive closed-form expressions for the new, key moments that we later target in the data.

Calibration. Armed with a new set of moments from the data, and a model built to interpret these moments, in Section 5 we calibrate the model to uncover the values of the structural parameters required to generate the patterns we observe in the data. This exercise produces both qualitative and quantitative insights into the effects of referrals on labor market outcomes.

First, interpreting the data through the lens of our model reveals the underlying relationships between a worker’s unobserved type, the frequency with which the worker meets firms through different job search channels, the quality of these potential matches, and the implications for labor market outcomes. We find that a significant amount of the heterogeneity in employment

and earnings is driven by the fact that some types of workers meet firms through business referrals much more frequently than other types. Referrals from family and friends, in contrast, are used more uniformly across worker types and generate relatively high productivity matches, on average, conditional on a worker’s type.

These relationships reveal a number of insights into the role that referrals play in the match formation process. The fact that business referrals are highly sensitive to a worker’s unobserved type—and, thus, the worker’s *ex ante* expected productivity—suggests that referrals from business contacts are used primarily to screen workers. Hence, we find that business referrals are best described by theories which ascribe a central role to a referrer’s ability to convey otherwise-private information, as in models based on adverse selection and homophily (e.g., Montgomery, 1991). Referrals from family and friends are most consistent with entirely different theories. Specifically, since referrals from family and friends tend to generate good matches independently of a worker’s underlying type, these referrals are better described by theories in which a referral improves match quality *ex post*—for example, by reducing symmetric uncertainty about idiosyncratic match quality (e.g., Simon and Warner, 1992) or easing inefficiencies that derive from moral hazard (e.g., Heath, 2018).

In addition to revealing qualitative insights into the role of referrals in the labor market, the calibrated model allows us to quantify the extent to which (different types of) referrals affect employment, earnings, inequality, and output across workers in high- and low-skill labor markets. Interestingly, though referrals from family and friends have a negative correlation with wages in our regression analysis, we find that they are a crucial source of jobs for a certain subset of workers that struggle to generate offers and matches through more traditional channels. For example, in the low-skill labor market, we find that denying a “low ability” worker referrals from family and friends would reduce the worker’s earnings by more than 11% and increase the worker’s likelihood of unemployment by 6 percentage points. Hence, despite concerns that referrals based on nepotism may exacerbate earnings inequality, our findings suggest that referrals from friends and relatives are, in fact, an important force for reducing earnings inequality.

Referrals from business contacts, in contrast, are used more frequently by high ability workers—who also receive offers through other channels at a high frequency—and their contribution to earnings is more pronounced in high skill occupations. Again, these results highlight the importance of interpreting our data through the lens of a structural model. The positive relationship between the use of a business referral and wages alone might have suggested that referrals from business contacts create high quality, productive matches. However, our model reveals that an important aspect of business referrals is that they increase the wages of workers who have relatively good employment prospects to begin with. Hence, the use of business referrals, which is typically encouraged by firms, increases output but also exacerbates earnings

inequality. These findings suggest that the implications of referrals may be more nuanced than they would appear at first glance, and speak to the ongoing debate regarding the sources of economic inequality and the design of policies aimed at mitigating its adverse effects.⁵

1.1 Related literature

Our paper contributes to the large (and growing) literature that studies the effects of referrals on labor market outcomes. For a broad overview of this literature, we refer the reader to the surveys by Ioannides and Loury (2004) and Topa (2011), and concentrate here on those studies most related to our work.

Most early attempts to quantify the impact of using a referral focused primarily on deriving empirical estimates of the productivity, wages, and turnover of referred workers, relative to non-referred workers (see, for instance, the seminal work by Datcher (1983), as well as Corcoran et al. (1980), Green et al. (1995), and Korenman and Turner (1996)). However, since referrals are typically not randomly assigned, these estimates are potentially driven by selection and unobserved heterogeneity, as opposed to capturing the direct effects of referrals. Broadly speaking, the literature has pursued two different approaches to overcome this challenge. First, a number of recent papers have exploited the availability of panel data and more sophisticated identification strategies to estimate the direct effects of referrals, including Bayer et al. (2008), Kramarz and Skans (2014), Schmutte (2015), Dustmann et al. (2016), Gee et al. (2017), and Heath (2018). Alternatively, a number of papers have used experimental settings to generate exogenous variation in the use of referrals; see, e.g., Bandiera et al. (2009), Beaman and Magruder (2012), Pallais and Sands (2016), and Friebe et al. (2023).

Our paper complements this strand of the literature in several important ways. The first derives from the unique nature of our data, which is drawn from a wide array of workers and occupations and contains detailed information about the job search process, including direct information about the use of different types of referrals. In contrast, most existing studies rely on data that either contains detailed information about job search methods *or* contains a representative sample of workers and/or occupations.⁶ Using these data, we document a new set of stylized facts that highlight the heterogeneous impacts of referrals, and provide a

⁵Topa (2019) and Hellerstein and Neumark (2020) discuss the implications of referrals for inequality.

⁶For example, the data used by Brown et al. (2016), Burks et al. (2015), Castilla (2005), Heath (2018), and Marmaros and Sacerdote (2002) contains direct evidence of whether a referral was used but are drawn from workers in specific occupations, industries, or demographic groups. Data collected from a more representative sample rarely contains information about how a worker-firm match was formed, forcing researchers to use proxies that are likely to be correlated with (specific types of) referrals: for example, Bayer et al. (2008) and Schmutte (2015) use geographic clustering, Dustmann et al. (2016) use ethnicity, Gee et al. (2017) use social media connections, Hensvik and Skans (2016) exploit overlap at a previous employer, and Kramarz and Skans (2014) use family relationships. However, unlike several of the datasets cited above, our data does not contain detailed information about firms, which limits our ability to control for firm fixed effects.

novel explanation for conflicting results in earlier studies. Second, we account for the role of unobserved worker heterogeneity and selection by interpreting the relationships we observe in the data through the lens of a model. Targeting a large number of moments from the data allows us to recover key model parameters, which reveal new qualitative insights into the role of (different types of) referrals in the matching process, along with quantitative estimates of the contribution of (different types of) referrals to employment, earnings, inequality, and output.

Our paper also complements the large theoretical literature that develops models to understand how referrals ease certain frictions in the matching process, and the subsequent implications for labor market outcomes. For example, some theories posit that referrals reduce adverse selection, since a current employee can provide information about a prospective worker’s unobserved productivity (Montgomery, 1991; Casella and Hanaki, 2008; Galenianos, 2014; Bolte et al., 2020). Other theories conjecture that referrals create good matches by reducing symmetric uncertainty regarding idiosyncratic match quality (Simon and Warner, 1992; Dustmann et al., 2016; Brown et al., 2016; Galenianos, 2013). Still others propose that referred workers are more productive *ex post* because the referrer can monitor or mentor the new worker (Kugler, 2003; Castilla, 2005; Beaman and Magruder, 2012; Heath, 2018). Finally, some models attribute the primary role of referrals to reducing search frictions by making workers better aware of existing vacancies (Holzer, 1988; Topa, 2001; Galeotti and Merlino, 2014; Galenianos, 2014; Schmutte, 2015).

In contrast to these papers, we use a model that does not derive specific microfoundations for a particular theory of referrals. Rather, we adopt a more flexible, parsimonious approach, and let the data dictate the relationship between, e.g., a worker’s underlying type, the rate at which the worker meets firms through different channels, the quality of the matches generated through each of these channels, and the subsequent effects on wages and turnover. By being *ex ante* agnostic about the specific mechanisms, the model remains rich enough to accommodate and identify the distinct properties of referrals from different sources. As a result, our paper sheds light on *which* of the theories cited above are most consistent with the empirical evidence regarding the relationship between different types of referrals and labor market outcomes across occupations.

Finally, our work is also closely related to several recent papers that study the impact of referrals on labor market outcomes by combining theoretical models of referrals with data containing detailed information on workers’ job search methods. Perhaps most closely related to our work is Arbex et al. (2019) and the contemporaneous paper by Moon (2023), both of which construct on-the-job search models and calibrate these models to data from the Survey of Consumer Expectations. Despite these broad similarities, the focus of these two papers is much different from our own. Arbex et al. (2019) develop a model with a rich network structure, and focus exclusively on heterogeneity in workers’ access to business referrals. Like us, they find

that this unobserved heterogeneity across workers is an important source of dispersion in employment status and earnings. However, their focus is more on the impact of network structure and connectedness, while we focus more on the distinct qualitative and quantitative effects of referrals from different sources across different occupations. Moon (2023) concentrates on developing explicit microfoundations for one particular theory of referrals—namely, he models the strategic incentives of a referrer to provide firms with an accurate signal of match-specific productivity. Consistent with our findings, his model predicts that referrals from business contacts should be associated with higher wages.⁷

2 Data

We use data from a supplement to the Survey of Consumer Expectations (SCE), which is administered by the Federal Reserve Bank of New York. The SCE is a nationally representative, monthly online survey of a rotating panel of about 1,300 household heads. New respondents are drawn each month to match various demographic targets from the American Community Survey (ACS), and they stay on the panel for up to twelve months. The supplement we use, called the Job Search Survey, has been administered annually since 2013.⁸

This dataset is particularly well-suited to our objectives in several dimensions. First, the survey asks a broad range of questions regarding how employed workers found their current job. Second, it asks about many different characteristics of the job including wages, benefits, job tenure, job satisfaction, and job search behavior. Third, since it is a representative survey, it covers workers across a wide range of individual characteristics and occupations. In addition to detailed information about the current job of the worker, our dataset also contains information about respondents’ previous work experience, along with all of the usual demographic information contained in the SCE data.

Our analysis focuses on non-self-employed individuals aged 18–64. This leaves us with a sample of about 5,000 observations covering the years 2013–2018. See Appendix A for additional details about how we generate some of our variables and arrive at our final estimation sample.

⁷Another recent, related paper is by Caldwell and Harmon (2019), who estimate a structural on-the-job search model using matched employer-employee data from Denmark. In contrast to our paper, they use a worker’s network of contacts as a source of variation in outside options to study the relationship between bargaining power and wages. They find that an increase in information about job openings—coming from closely connected former co-workers employed in other firms—leads to higher mobility and wage growth.

⁸The survey was designed by Jason Faberman, Andreas Mueller, Ayşegül Şahin, and Giorgio Topa. See Faberman et al. (2022) for a more detailed description of the survey and associated dataset.

2.1 Construction of key variables

Before presenting our empirical results, we describe how we construct two key variables related to the source of the referral (business versus family/friends) and the skill content of the job.

First, to determine whether a worker used a referral, and the *type of the referral* used, we rely on a question from the survey that asks currently employed workers how they “learned about their current job”. Using the worker’s response to this question, we construct binary indicators for two types of referrals: (i) family or friend, and (ii) business contact. Since individuals are allowed to give multiple responses to this question, these measures are not mutually exclusive. For those that indicated they were “referred by a friend or relative,” we set the indicator for referral from family and friends equal to 1, and we set it to 0 otherwise. For referral from business contacts, we set the indicator equal to 1 if the individual responded that they were “referred by a former co-worker, supervisor, business associate”. We also set the business contacts indicator equal to 1 if they reported being “referred by a current employee at the company,” as long as they did not also indicate that they were referred by a friend or relative. In other words, if a worker who indicated that they were referred by a friend or relative also indicated that they were referred by a current employee at the company, we classify this as a referral from a family member or friend, as it seems most likely that the two answers correspond to the same referrer. However, if the worker responded that they were referred by a current employee at the firm but *not* by a friend or relative, we classify the referrer as a business contact.⁹

Second, to classify *different types of jobs*, we measure the skill content of each (employed) worker’s reported occupation using the Nam-Powers-Boyd (NPB) occupational index. This index ranks occupations (at the 3-digit occupation level) based on the earnings and educational levels of the workers in each occupation.¹⁰ To do so, one first calculates the median education level and median earnings of individuals in each occupation. Then, these values are weighted by the number of people in each occupation to create a percentile measure of the position of each occupation in both the education and earnings distributions. Finally, these two percentiles are averaged to generate the index.¹¹ The version we use comes from 2016 and is based on data from the American Community Surveys from 2010-2012, accessible via IPUMS.¹²

⁹Using this method of classification, about 8% of referred workers have both referrals indicators equal to 1. We experimented with several ways of dealing with the overlap between the two measures, including fully partitioning the three responses related to referrals, and our empirical results did not change significantly.

¹⁰Occupations in the SCE are categorized using the Standard Occupational Classification System (SOC) from the Bureau of Labor Statistics (BLS).

¹¹We also experimented with an alternative occupation index constructed using O*NET data. Specifically, the measure was computed as the fraction of jobs within an occupation code that require a bachelor’s degree, which generated scores that also ranged from 0 to 100. Results were qualitatively similar using this alternative measure.

¹²The NPB scores are available for download at <http://www.npb-ses.info/>.

To give the reader a sense of the NPB occupational index, Table 9 in Appendix A provides a list of NPB scores assigned to various occupations, aggregated at the 2-digit occupation level for the sake of presentation. Scores range from 0 to 100, with “Food Preparation and Serving Related Occupations” (FOOD) at the bottom and “Legal Occupations” (LEGL) at the top. Note that each of these groups is a weighted average of scores at the 3-digit occupation level; for example, FOOD contains both “chefs and head cooks” (NPB score of 40) and “dishwashers” (NPB score of 1), while LEGL contains both “lawyers, judges, and related workers” (NPB score of 99) and “paralegals and legal assistants” (NPB score of 70). For all of our regression analysis below, we use the finer, 3-digit occupation scores.

3 Empirical Results

In this section, we examine the frequency with which the two different types of referrals are used across occupations, and the characteristics of matches formed using each type of referral. Since our indicators of usage are derived from currently employed workers, our estimates in this section are based on the sample of workers who were currently employed at the time of the survey. Note that in Section 5, when we calibrate our model to a larger set of moments, we will use the full sample of employed and unemployed workers.

3.1 Usage of Referrals Across Occupations

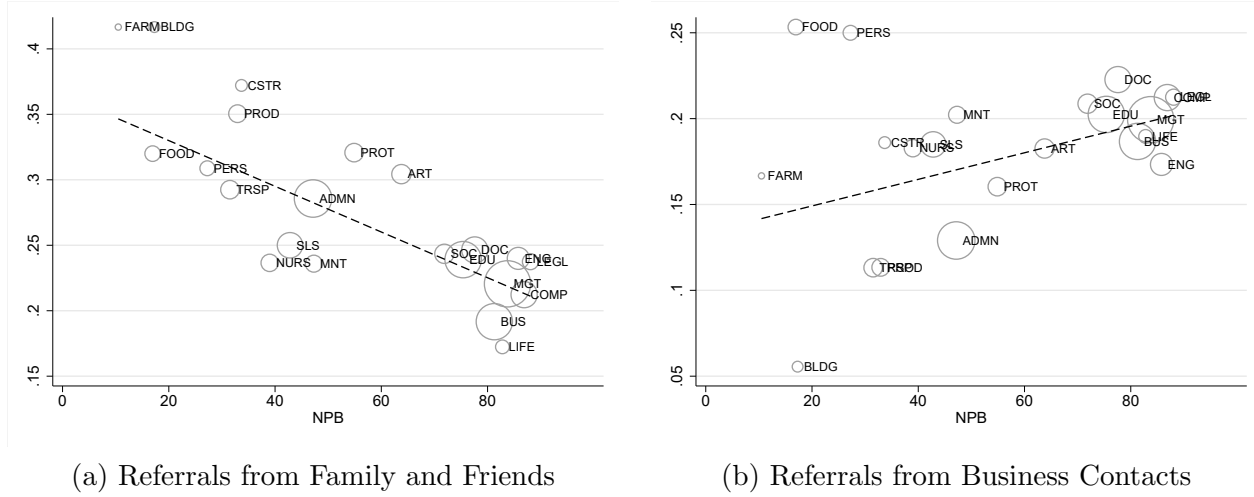
We first examine the relationship between the usage of the two types of referrals and the skill requirements of different occupations. As a first step, Figure 1a plots, for each 2-digit occupation code, the percentage of currently employed workers who report having used a referral from a family member or friend in the process of being hired at their current job. Figure 1b plots the corresponding relationship for business referrals. The figures suggest that referrals from family and friends are used more often in the formation of low-skill jobs, while referrals from business contacts are used relatively more often in the formation of high-skill jobs.

Of course, these patterns could reflect differences in the characteristics of the workers in these occupations, and not necessarily differences in the occupations themselves. To establish that this relationship is not just capturing worker characteristics, we run a linear regression on a dummy variable for referral usage (for each type of referral) against the skill index of the occupation, time and geographic region fixed effects, and a rich set of worker characteristics.^{13,14} These characteristics include age, gender, race, marital status, number of children under the

¹³While we employ a linear probability model for all of our binary outcomes, for ease of exposition, our results are very similar using a logit or probit specification.

¹⁴In the regressions, we use the skill index for the more detailed 3-digit occupation code. However, the results remain similar if we use the more aggregated 2-digit occupation code, as in Figures 1a and 1b above.

Figure 1: The Use of Referrals Across Occupations



Notes: This figure plots the fraction of individuals within each 2-digit occupation that found their current job through a referral from family and friends (a) and from a business contact (b) against the skill content of the job (NPB score). The size of each dot is proportional to the number of individuals within each occupation.

age of 6, and home ownership status.¹⁵ Table 1 confirms that there is a positive relationship between the use of business referrals and occupational skill, and a (stronger) negative relationship between the use of referrals from family and friends and occupational skill. To interpret these results, comparing a worker in “Retail Sales” (NPB score of about 26) with a worker in “Media and Communication” (NPB score of about 76), our results imply that the “Media and Communication” worker is 3pp (18%) more likely to use a business referral and 10pp (39%) less likely to use a family/friends referral.

The patterns we document in Figure 1 (and Table 1) suggest that referrals from family and friends and referrals from business contacts might be playing different roles (or helping to overcome different frictions) in the matching process. This observation prompts two conjectures. First, if the two types of referrals are playing different roles, then one might naturally expect them to be associated with different labor market outcomes. Second, if the mix of referrals from family and friends and business contacts varies across occupations—and referrals from these two sources have different effects on labor market outcomes—then one would also expect that two studies focusing on different occupations or sectors may find conflicting results regarding the relationship between labor market outcomes and the use of *any* type of referral.

We now explore these conjectures in the data. In particular, we examine the relationship between the use of referrals and two standard measures of labor market outcomes: wages and

¹⁵While we observe detailed information in our data about workers, we observe little about employers outside of their industry and measures of firm size. However, all of the empirical results in the paper are robust to controlling for both industry and firm size.

Table 1: Referral Usage and Skill Index

	Type of Referral			
	Business	Family/Friends	Business	Family/Friends
Skill Index	0.0008*** (0.0003)	-0.0018*** (0.0003)	0.0006** (0.0003)	-0.0020*** (0.0003)
Time and Region FE			✓	✓
Individual Controls			✓	✓
N	3779	3779	3779	3779

Notes: Estimates are from regressions in which the outcome is whether an individual used either a business or family/friend referral to find their current job. Individual controls include age, gender, race, marital status, number of children under the age of 6, and home ownership status. Standard errors are presented in parentheses below the point estimates. *** denotes significance at the 1% level, ** denotes significance at the 5% level, and * denotes significance at the 10% level.

tenure. We show that clear and *opposing* relationships emerge, but only after conditioning on the two types of heterogeneity highlighted above—namely, different types of referrals and different types of occupations.

3.2 Referrals and Starting Wages

We first study workers’ starting wages. In column (1) of Table 2, we report the results of a regression of log (real) starting wages on dummy variables that indicate whether the worker used a referral from a business contact or family/friend in the hiring process. Again, we control for time and region fixed effects, as well as observable worker characteristics. We find that workers referred to their current job by a business contact have starting wages that are approximately 16% higher than those of non-referred workers, while those referred by family and friends have starting wages that are approximately 9% lower than those of non-referred workers. In column (2), we control for the skill index of the worker’s occupation.¹⁶ The coefficient on business referrals is essentially unchanged, while the coefficient on referrals from family and friends decreases in absolute value, but remains negative and statistically significant.

In column (3), we also control for the previous wage in an attempt to control for unobserved worker heterogeneity. Not surprisingly, there is a strong positive relationship between the wage at the previous job and the starting wage at the current job. Moreover, while the coefficient on business referrals remains positive and statistically significant, the coefficient on referrals from family and friends becomes small and insignificant. As we discuss in more detail below, these findings are consistent with a selection effect; that is, referrals from family and friends do not

¹⁶Overall our regression results are similar if, instead of conditioning on the skill index measure of occupations (NPB score), we use 3-digit occupation dummy variables.

Table 2: Starting Wages and Referrals

	Log Real Starting Wage					
	(1)	(2)	(3)	(4)	(5)	(6)
Business Referral	0.161*** (0.025)	0.148*** (0.023)	0.085*** (0.022)			
Family/Friends Referral	-0.093*** (0.023)	-0.046** (0.021)	-0.024 (0.020)			
Any Referral				0.005 (0.020)	0.028 (0.019)	0.017 (0.017)
Skill Index		0.010*** (0.000)	0.005*** (0.000)		0.010*** (0.000)	0.005*** (0.000)
Log Previous Wage			0.530*** (0.014)			0.535*** (0.014)
Time and Region FE	✓	✓	✓	✓	✓	✓
Individual Controls	✓	✓	✓	✓	✓	✓
N	3317	3317	2311	3317	3317	2311

Notes: Estimates are from regressions of the log of the real starting wage for the worker's current job. Individual controls include age, gender, race, marital status, number of children under the age of 6, and home ownership status. Standard errors are presented in parentheses below the point estimates. *** denotes significance at the 1% level, ** denotes significance at the 5% level, and * denotes significance at the 10% level. There are 462 observations for which we do not observe the starting wage. We lose 1006 observations when adding previous wage as a control due to missing data on previous wage. Results for the specifications without previous wage as a control (columns 1, 2, 4, and 5) are similar when using this more restricted sample.

necessarily cause lower wages, but rather these types of referrals tend to be used by workers who earn lower starting wages because of characteristics not easily observed by the econometrician. For example, the results in column (3) are consistent with an environment in which workers who use referrals from family and friends have fewer outside options than otherwise similar workers, and hence receive lower wages.

Columns (4)–(6) report results for the same regressions without distinguishing between the two types of referrals, i.e., we regress starting wages on a dummy variable that takes value 1 if the worker used *any* type of referral. The relationship between the use of a referral and starting wages disappears. This insight—that referrals from business contacts and family/friends are associated with opposing, offsetting effects on starting wages—may explain why the existing literature has found mixed evidence regarding the relationship between wages and referrals.

3.3 Referrals and Job Tenure

We now analyze the relationship between the use of referrals and job tenure. We measure the job tenure of all currently employed workers at the time of the survey using data on the start

date of each worker’s current job. Ideally, one would like to analyze the duration of completed employment spells, but this is not possible given the repeated, cross-sectional nature of our data.¹⁷ However, despite the potential limitations of our stock-sampled measure of job tenure, the results based on this measure reveal crucial information about the relative tenure of workers across job search methods and occupations. For one, it is well known that data of this nature—which is left-truncated and right-censored—suffers from competing biases that, under certain assumptions, cancel each other out.¹⁸ Moreover, even if they do not offset each other exactly, our focus on *relative tenure*—and the large differences we find between the relative tenure of workers who used business referrals and those who used referrals from family and friends—would likely dominate any biases due to stock sampling. Finally, as a robustness check, we use our quantitative model to directly compare stock-sampled spells to completed spells, and we find that the two measures deliver almost identical results regarding the differences in tenure across job search method.

In Table 3, we regress our measure of job tenure on the dummy variables for referrals, using the same set of controls described above. Columns (1)–(3) show that workers who were referred by business contacts have significantly shorter job durations than non-referred workers, while those who were referred by family and friends have significantly longer durations. Note that the result for business referrals would be considered surprising within the context of theories with symmetric uncertainty regarding match-specific productivity (à la Jovanovic, 1979), since such theories typically predict a positive relationship between match quality, wages, and job tenure. Indeed, columns (4)–(6) suggest that one would actually find a positive (although not precisely estimated) relationship between the use of any referral and job tenure in our data. However, much like our results on wages and referral usage, these regressions would be misleading, since they mask stark differences between the effects of referrals from business contacts and those from family and friends.

What generates these patterns? There are many reasons why matches that are formed through different channels may last for longer or shorter periods of time. For example, if relatives and friends have superior knowledge about a worker’s preferences or personal circumstances, then matches formed through family and friends could be “better” along non-pecuniary dimensions, such as flexible hours, non-wage benefits, or the potential for faster advancement, which would explain why workers who match through this channel tend to stay at their job

¹⁷Since workers are only tracked for one year in the SCE, we observe very few completed job spells.

¹⁸On the one hand, since our sample is left-truncated, workers with shorter spells are less likely to be sampled, which leads to overestimating the average length of employment spells. On the other hand, since it is right-censored, our measure of average tenure underestimates the average length of completed spells. Under certain assumptions (see, e.g., Heckman and Singer, 1984), these biases cancel and the average job duration that we observe is a consistent estimate of the true, uncensored duration.

Table 3: Job Tenure and Referrals

	Log Job Duration					
	(1)	(2)	(3)	(4)	(5)	(6)
Business Referral	-0.185*** (0.055)	-0.195*** (0.054)	-0.172*** (0.065)			
Family/Friends Referral	0.216*** (0.049)	0.246*** (0.049)	0.229*** (0.060)			
Any Referral				0.070 (0.043)	0.084* (0.043)	0.087* (0.052)
Skill Index		0.007*** (0.001)	0.007*** (0.001)		0.006*** (0.001)	0.007*** (0.001)
Log Previous Wage			-0.142*** (0.036)			-0.157*** (0.036)
Time and Region FE	✓	✓	✓	✓	✓	✓
Individual Controls	✓	✓	✓	✓	✓	✓
N	3779	3779	2476	3779	3779	2476

Notes: Estimates are from regressions of the log of the duration of the current job. Individual controls include age, gender, race, marital status, number of children under the age of 6, and home ownership status. Standard errors are presented in parentheses below the point estimates. *** denotes significance at the 1% level, ** denotes significance at the 5% level, and * denotes significance at the 10% level. We lose 1303 observations when adding previous wage as a control due to missing data on previous wage. Results for the specifications without previous wage as a control (columns 1 and 2) are similar when using this more restricted sample.

longer. However, we do not find any evidence suggesting a relationship between the use of referrals and job satisfaction or “fit.” In particular, in Appendix A, we exploit several questions from the survey on job satisfaction to document that workers hired through either type of referral are no more or less satisfied with various aspects of their job than non-referred workers.¹⁹ Consistent with this evidence, we also show that workers hired through either type of referral are no more or less likely to be currently looking for a new job, relative to non-referred workers. We also examined wage growth and found no significant differences across job-finding method.

Instead, the difference in job tenure appears to be driven by different arrival rates of outside offers *after* being hired into their current job. Table 4 reports the output of a linear regression model where the dependent variable is an indicator for whether a currently employed worker has had at least one contact with another firm in the last four weeks. As is evident, workers who got their current job through a business contact are significantly more likely to make contact with additional firms than non-referred workers, whereas those who were hired through a referral from a family member or friend are significantly less likely to have generated new contacts in

¹⁹As we discuss in detail in the Appendix, workers are asked about their overall satisfaction, their satisfaction with their compensation, the “fit” of the job, the opportunities for promotions or other forms of career progression, and their satisfaction with other, non-wage aspects of the job.

the past four weeks.²⁰

Table 4: Contact Rates and Referrals

	Probability of Contact (Last 4 Weeks)		
	(1)	(2)	(3)
Business Referral	0.050*** (0.018)	0.048*** (0.018)	0.037* (0.023)
Family/Friends Referral	-0.043*** (0.016)	-0.037** (0.016)	-0.050** (0.021)
Skill Index		0.001*** (0.000)	0.001 (0.000)
Log Previous Wage			0.047*** (0.013)
Time and Region FE	✓	✓	✓
Individual Controls	✓	✓	✓
N	3779	3779	2476

Notes: Estimates are from regressions of an indicator for whether or not an individual had contact with at least one potential employer in the last four weeks. Individual controls include age, gender, race, marital status, number of children under the age of 6, and home ownership status. Standard errors are presented in parentheses below the point estimates. *** denotes significance at the 1% level, ** denotes significance at the 5% level, and * denotes significance at the 10% level. We lose 1303 observations when adding previous wage as a control due to missing data on previous wage. Results for the specifications without previous wage as a control (columns 1 and 2) are similar when using this more restricted sample.

3.4 Summary of Facts and Key Ingredients for a Model

In this section, we established that the frequency with which different job search channels are used varies systematically across occupations, and that worker-firm matches formed through different job search channels are associated with significant differences in labor market outcomes. In particular, business referrals are used relatively more frequently at high skill jobs and are associated with higher starting wages but shorter tenures, as workers hired through business referrals continue to receive offers at a high rate after forming a match. In contrast, matches formed with the help of family and friends occur relatively more frequently at low skill jobs and are associated with lower starting wages, though this difference vanishes when controlling for unobserved worker heterogeneity (through previous wages). Still, despite earning relatively low wages, workers who are hired through family and friends tend to stay longer at their job, as they receive new opportunities less frequently than others.

²⁰We test other horizons as well and find similar results.

In the next section, we use these facts to guide the construction of a structural model of the labor market. The model has three key ingredients. First, we assume that matches between workers and firms can be formed through *three distinct channels*: a referral from family and friends, a referral from a business contact, or other (formal) job search channels. Second, since workers in our data generate offers through these different channels at (persistently) different rates, even after controlling for a variety of observable characteristics, we introduce *unobserved worker heterogeneity*. Importantly, we allow the rate at which workers meet firms through the different job search channels and the quality of the matches they form to depend on their unobserved type or “ability.” Third, since workers’ (heterogeneous) abilities to generate offers do not vanish after forming a match, we assume that workers search both off and *on the job*.

However, one might ask, why do we need a model at all? We argue that there are (at least) two important reasons to interpret our empirical findings through the lens of a structural model. First, while any number of stories could qualitatively fit a subset of the empirical patterns we document, an equilibrium model allows us to identify and test various mechanisms across multiple, inter-related dimensions. Consider, for example, the positive relationship between the use of a business referral and a worker’s starting wage. Studying this bivariate relationship alone, one might be tempted to conclude that matches formed through business referrals are more productive, perhaps because referrals provide an efficient technology for sharing information about (*ex post*) idiosyncratic match quality. However, an equally plausible explanation is that (*ex ante*) more productive workers tend to use business referrals to find a job, perhaps because reputation concerns ensure that business contacts only refer “good” workers. Yet another possibility is that matches formed through business referrals are no more productive than matches formed through other channels, but instead workers who tend to use business referrals have better outside options (perhaps because they have larger networks of business contacts) and hence negotiate higher wages.

Distinguishing between these candidate explanations is crucial for understanding the role of business referrals in the job-finding process and the subsequent implications for labor market outcomes, at both the individual level and in the aggregate. An equilibrium model imposes a certain amount of discipline—and hence offers a more stringent test—for any potential explanation of our empirical findings. Continuing with the example above, while it may be straightforward to identify a mechanism that generates a wage premium for workers hired through a business referral, it is considerably more challenging to require that the candidate mechanism also explains why these workers leave their jobs more quickly, why this channel is used more frequently at high skill jobs, why the wage premium associated with business referrals is more pronounced at high skill jobs, and why referrals from family and friends generate entirely different patterns altogether.

4 Model

We incorporate multiple job search channels into the workhorse model of on-the-job search with unobserved worker heterogeneity, as formulated by Postel-Vinay and Robin (2002) and Cahuc et al. (2006). In fact, as we establish below, if one aggregates these multiple channels into a single matching technology, the characterization of many standard equilibrium objects is essentially the same as in Cahuc et al. (2006). For this reason, we keep the characterization of these objects brief (though sufficiently detailed to remain self-contained), and instead focus on deriving the novel predictions that emerge from our framework—namely, the endogenously-generated relationships between the channel through which a worker found her current job, the arrival rate and quality of new job opportunities, and the subsequent implications for her wage and tenure. In particular, we are able to derive analytical expressions from our model corresponding to all key moments from the data (channel-specific and otherwise). Hence, despite a seemingly minimal departure from a well-known benchmark, our model is rich enough to confront the new facts that we uncovered in Section 3, with essentially no loss in tractability.

4.1 Environment

We consider a continuous time, infinite horizon environment. All agents are risk neutral and discount the future at rate $r > 0$. There is a measure 1 of workers who are heterogeneous with respect to an unobserved attribute or “ability” $a \in \mathcal{A} \equiv \{a_1, \dots, a_N\}$ for some $N \in \mathbb{N}$. We let π_i denote the fraction of workers with ability a_i , with $\sum_{i=1}^N \pi_i = 1$. As we discuss below, a worker’s ability a_i should not be confused with their skill or occupation: when we take the model to the data, we interpret each occupation or skill level as a separate labor market, and interpret a_i as the unobservable ability of workers *within that market*.²¹

There is a large measure of firms that operate a constant returns-to-scale production technology. When a worker meets a firm, the pair draws a match-specific productivity $x \in [\underline{x}, \bar{x}]$. If they choose to form a match, the worker and firm jointly produce a flow amount $f(x) = px + c$ for some $p \in \mathbb{R}_+$ and $c \in \mathbb{R}$. An unmatched (unemployed) worker consumes a flow amount b , while an unmatched vacancy at a firm produces 0. Worker-firm matches are exogenously destroyed at rate δ .

Meetings. The first departure from the existing literature is that we assume contacts or “meetings” between workers and firms occur through one of three channels: a referral from a family member or friend; a referral from a business contact; or formal (“other”) channels.

²¹For example, our model could be interpreted as the market for lawyers, and a_i distinguishes the unobserved ability *across* lawyers.

We denote these by F , B , and O , respectively, and denote the set of possible channels by $\mathcal{C} \equiv \{F, B, O\}$.

The second departure from the existing literature is that we assume a worker's type can affect the rate at which he meets firms through the various channels. In particular, we assume that employed and unemployed workers of ability a_i generate meetings through channel $j \in \mathcal{C}$ at rate $\lambda_j^e(a_i)$ and $\lambda_j^u(a_i)$, respectively. Conditional on meeting, a match-specific productivity is then drawn from a distribution with cdf $H_j(x|a_i)$. It will be convenient to define

$$\Gamma_j^k(x|a_i) = \lambda_j^k(a_i) \bar{H}_j(x|a_i)$$

for $j \in \mathcal{C}$ and $k \in \{e, u\}$, where $\bar{H}_j(x|a_i) \equiv 1 - H_j(x|a_i)$. In words, $\Gamma_j^e(x|a_i)$ is the arrival rate of offers for an employed worker of ability a_i through channel j with a match-specific productivity that exceeds x . It will also be convenient to define the arrival rate of such contacts through *any* channel by

$$\Gamma^k(x|a_i) = \sum_{j \in \mathcal{C}} \Gamma_j^k(x|a_i), \quad k \in \{e, u\}.$$

This specification allows a worker's type to affect both the arrival rate of meetings and the idiosyncratic quality of the match. This modeling choice, while stylized, is meant to encapsulate a variety of micro-founded theories of job referrals (discussed in the literature review), as some theories focus on the role of referrals in generating meetings for (at least some types of) workers, while other theories focus on the quality of matches formed through referrals.²²

Wage Determination. To close the model, we assume that wages are determined by the strategic wage-bargaining protocol described in Cahuc et al. (2006). According to this protocol, when an unemployed worker meets a firm and there are gains from trade, the firm and the worker bargain over the wage as in standard models (e.g., Mortensen and Pissarides, 1994). We let β denote the share of the surplus that the worker receives, or the worker's "bargaining power".

When an employed worker meets a new firm, a three-player bargaining game ensues. If the match-specific productivity at the poaching firm (x') is greater than at the incumbent firm (x), the worker moves to the poaching firm. The worker and the poaching firm bargain over the wage, where we define the worker's outside option (of not moving to the poaching firm) as remaining employed at the incumbent firm at a wage equal to his marginal productivity, $f(x)$, which is the maximum that the incumbent firm would agree to pay him.

If $x' < x$, however, the worker remains at the incumbent firm, but his wage might be adjusted. In particular, if the expected value of remaining at the incumbent firm at wage w

²²The distinction between referrals' effect on meetings and match quality is also present in a recent working paper by San (2021).

is less than the outside option of moving to the poaching firm at the maximum wage $f(x')$, the worker remains at the incumbent firm but renegotiates his wage using the outside option of the poaching firm. Otherwise, the worker remains at the incumbent firm and his wage is unchanged.

4.2 Key Equilibrium Objects

In this section, we define the key objects that make up a steady-state equilibrium in the model described above. Proposition 1, below, provides a complete characterization. As the derivation is relatively close to that in Cahuc et al. (2006), it is relegated to Appendix C.

Value Functions and Wage Functions Let $V^u(a_i)$ denote the expected discounted value of an unemployed worker with ability a_i , and let $V^e(a_i, x, w)$ denote the expected discounted value of a worker with ability a_i who is currently employed at a firm with match-specific productivity x earning a wage w . Since a firm generates zero output when unmatched, the expected surplus created by forming a match with productivity x is $V^e(a_i, x, f(x)) - V^u(a_i)$, i.e., the worker's value of being employed at a wage equal to the total output of the match, $f(x)$, less the worker's value of being unemployed. Hence, it is straightforward to establish that an unemployed worker with ability a_i will form a match with a new firm if, and only if, the match-specific productivity $x \geq x^*(a_i) \equiv x_i^*$, where x_i^* satisfies

$$V^u(a_i) = V^e(a_i, x_i^*, f(x_i^*)), \quad i \in \{1, \dots, N\}. \quad (1)$$

Following Cahuc et al. (2006), the worker will earn a wage $w^u(a_i, x)$ that satisfies

$$V^e(a_i, x, w^u(a_i, x)) = V^u(a_i) + \beta [V^e(a_i, x, f(x)) - V^u(a_i)]. \quad (2)$$

Intuitively, $w^u(a_i, x)$ yields the worker an expected utility equal to his outside option of unemployment plus a share β of the match surplus.

Now consider an employed worker with ability a_i , productivity x , and wage w who contacts a new firm and draws match-specific productivity x' . If $x' > x$, the worker moves to the new firm at a wage $w^e(a_i, x, x')$ satisfying

$$V^e(a_i, x', w^e(a_i, x, x')) = V^e(a_i, x, f(x)) + \beta [V^e(a_i, x', f(x')) - V^e(a_i, x, f(x))]. \quad (3)$$

Intuitively, $w^e(a_i, x, x')$ yields the worker an expected utility equal to his outside option of remaining at the incumbent firm at the highest wage they are willing to pay, $V^e(a_i, x, f(x))$, plus a share β of the surplus created by moving to the more productive match, $V^e(a_i, x', f(x')) -$

$V^e(a_i, x, f(x))$.

Alternatively, if $x' \leq x$, the worker will remain at his current job, though he will use the threat of leaving to renegotiate his current wage if x is sufficiently high. In particular, let $\hat{x}(a, x, w)$ denote the value of x' such that a currently employed worker of ability a with match-specific productivity $x > x'$ would renegotiate her wage to exactly w , i.e.,

$$w = w^e(a_i, \hat{x}(a, x, w), x), \quad (4)$$

so that $w^e(a_i, x', x) \leq w$ if $x' \leq \hat{x}(a, x, w)$. Then the worker renegotiates her wage to $w(a_i, x', x)$ if $x' \geq \hat{x}(a_i, x, w)$, and otherwise her wage remains w .

Distribution of Workers Unemployed workers of ability a_i exit unemployment when they meet a firm and draw match-specific productivity $x \geq x_i^*$. Once employed, a worker moves only when he meets a new firm with a higher match-specific productivity. In this section, we use these transition rules to derive the distribution of workers across possible states.

To do so, let $\phi^u(a_i)$ denote the measure of unemployed workers with ability a_i , and let $\phi^e(a_i, x)$ denote the measure of workers with ability a_i currently employed at a job with match-specific productivity x . It will be convenient to define the *cumulative measure* of employed workers with ability a_i and match-specific productivity $x' \leq x$ by $\Phi^e(x|a_i) \equiv \int_{\underline{x}}^x \phi^e(a_i, x') dx'$. These equilibrium objects are characterized by three sets of conditions:

$$\pi_i = \phi^u(a_i) + \Phi^e(\bar{x}|a_i) \quad (5)$$

$$\dot{\phi}^u(a_i) = \delta \Phi^e(\bar{x}|a_i) - \phi^u(a_i) \Gamma^u(x_i^*|a_i) = 0 \quad (6)$$

$$\dot{\Phi}^e(x|a_i) = -\Phi^e(x|a_i) [\delta + \Gamma^e(x|a_i)] + \phi^u(a_i) [\Gamma^u(x_i^*|a_i) - \Gamma^u(x|a_i)] = 0 \quad (7)$$

for all $i \in \{1, \dots, N\}$ and $x \in [x_i^*, \bar{x}]$. Condition (5) simply requires that summing the measures of unemployed and employed workers with ability a_i yields the (exogenously specified) aggregate measure of workers with ability a_i . Conditions (6) and (7) are steady-state conditions, equating the inflow and outflow of workers into every possible state.²³ For example, in equation (7), workers with ability a_i and current productivity $x' \leq x$ exit when their match is destroyed, which occurs at rate δ , or when they find a better match with productivity $x'' > x$, which occurs at rate $\Gamma^e(x|a_i)$. Meanwhile, unemployed workers with ability a_i enter this state by meeting a firm and drawing productivity $x' \in [x_i^*, x]$, which occurs at rate $\Gamma^u(x_i^*|a_i) - \Gamma^u(x|a_i)$.

²³We use the notation of, e.g., $\dot{\phi}^u(a_i)$ to denote the rate of change in the measure of unemployed workers of type a_i .

4.3 Equilibrium

A steady-state equilibrium is characterized by thresholds $x^*(a)$ and $\hat{x}(a, x, w)$, value functions $V^u(a)$ and $V^e(a, x, w)$, wage functions $w^u(a, x)$ and $w^e(a, x, x')$, and distribution functions $\phi^u(a)$ and $\Phi^e(x|a)$ satisfying equations (1)–(7).²⁴ The following proposition provides a closed-form characterization of the equilibrium objects that are key to our analysis.²⁵

Proposition 1. *In a steady-state equilibrium, the wage functions are given by*

$$w^e(a_i, x, x') = f(x') - p(1 - \beta) \int_x^{x'} \frac{r + \delta + \Gamma^e(x''|a_i)}{r + \delta + \beta\Gamma^e(x''|a_i)} dx'', \quad (8)$$

for $x' > x \geq x_i^*$, and

$$w^u(a_i, x) = w^e(a_i, x_i^*, x), \quad (9)$$

for $x \geq x_i^*$, where x_i^* satisfies

$$f(x_i^*) = b + p\beta \int_{x_i^*}^{\bar{x}} \frac{[\Gamma^u(x|a_i) - \Gamma^e(x|a_i)]}{r + \delta + \beta\Gamma^e(x|a_i)} dx \quad (10)$$

for $i \in \{1, 2, \dots, N\}$. The distribution functions are given by

$$\phi^u(a_i) = \frac{\delta\pi_i}{\delta + \Gamma^u(x_i^*|a_i)} \quad \text{for all } i \in \{1, \dots, N\}, \quad (11)$$

and

$$\Phi^e(x|a_i) = \frac{\delta\pi_i [\Gamma^u(x_i^*|a_i) - \Gamma^u(x|a_i)]}{[\delta + \Gamma^u(x_i^*|a_i)] [\delta + \Gamma^e(x|a_i)]} \quad \text{for all } x \geq x_i^* \text{ and } i \in \{1, \dots, N\}. \quad (12)$$

In what follows, it will be helpful to derive the (unconditional) distribution of wages across workers, which we do in the Lemma below.

Lemma 1. *For any wage $w \in [w^u(a_i, x), w^e(a_i, x, x)]$, the fraction of workers of type a_i employed at a firm with productivity $x \geq x_i^*$ that earn a wage $w' \leq w$ is given by*

$$G(w|a_i, x) = - \frac{\phi^u(a_i) d\Gamma^u(x|a_i) + \Phi^e(\hat{x}(a_i, x, w)|a_i) d\Gamma^e(x|a_i)}{\phi^e(a_i, x) [\delta + \Gamma^e(\hat{x}(a_i, x, w)|a_i)]}. \quad (13)$$

4.4 Key Model-Implied Moments

We now show how the framework developed above can be used to analyze and interpret the empirical regularities documented in Section 3. As a first step, we derive the joint distribution

²⁴Recall that, for ease of exposition, we denote $x^*(a_i)$ by x_i^* .

²⁵Given these expressions, the remaining equilibrium objects can be easily constructed. The proof is in Appendix C.

of employed workers' unobserved types, their match-specific productivities, and the channel through which they found their job. Then we use this distribution to derive analytical expressions for several key model-implied moments, including the fraction of workers that used each job search channel, and the relationship between a worker's job search channel, wage, and expected tenure. In Section 5, we use these expressions—which, to the best of our knowledge, are new to the literature—to derive quantitative estimates of the model's structural parameters from key moments in the data, without resorting to costly simulations.

As a first step, since workers' future labor market transitions do not depend on the channel through which they formed their current match, the probability that a worker of type a_i currently employed with productivity x received her job through channel $j \in \{B, F, O\}$ is

$$\Lambda_j(a_i, x) = \frac{\phi^u(a_i)d\Gamma_j^u(x|a_i) + \Phi^e(x|a_i)d\Gamma_j^e(x|a_i)}{\sum_{j \in \{B, F, O\}} \phi^u(a_i)d\Gamma_j^u(x|a_i) + \Phi^e(x|a_i)d\Gamma_j^e(x|a_i)}.$$

To see why, note that the numerator represents the flow of (unemployed and employed) type a_i workers into matches of quality x through channel j , while the denominator represents the flow of type a_i workers into matches of quality x through any channel.

Using this expression, we can define the measure of workers of type a_i currently employed with productivity x that got their job through channel j as

$$\phi_j^e(a_i, x) = \Lambda_j(a_i, x)\phi^e(a_i, x). \quad (14)$$

Integrating and summing $\phi_j^e(a_i, x)$ reveals that the fraction of currently employed workers who used channel j to find their current job is

$$\frac{1}{1-u} \sum_i \int_{x_i^*}^{\bar{x}} \phi_j^e(a_i, x) dx, \quad (15)$$

where $u = \sum_i \phi^u(a_i)$ denotes the measure of unemployed workers.²⁶

Since the distribution of wages across workers, conditional on a_i and current productivity x , is the same for all $j \in \{B, F, O\}$,²⁷ the average wage of currently employed workers who got

²⁶Since the measure of workers is normalized to one, note that u is also equal to the unemployment rate.

²⁷To see why, note that the current wage of a type a_i worker with productivity x , $w^e(a_i, x', x)$, only depends on the value $x' < x$ of either his last job or his last offer (with $x' = x^*$ if he was last unemployed). Given the nature of Poisson arrivals, x' does not depend on the channel through which the worker got his current job.

their job through channel j is thus²⁸

$$\frac{\sum_i \int_{x_i^*} \mathbb{E}[w | a_i, x] \phi_j^e(a_i, x) dx}{\sum_i \int_{x_i^*} \phi_j^e(a_i, x) dx}, \quad (16)$$

where, letting $\underline{w} \equiv w^u(a_i, x)$ and $\bar{w} \equiv w^e(a_i, x, x)$,

$$\mathbb{E}[w | a_i, x] = \underline{w} G(\underline{w} | a_i, x) + \int_{\underline{w}}^{\bar{w}} w dG(w | a_i, x) = \bar{w}(a_i, x) - \int_{\underline{w}}^{\bar{w}} G(w | a_i, x) dw. \quad (17)$$

Similarly, since the expected tenure of a worker of type a_i who is currently employed at a firm with productivity x ,

$$\tau(x, a_i) = \frac{1}{\delta + \Gamma^e(x | a_i)},$$

is also independent of the channel through which the worker got the job, the expected tenure of currently employed workers who got their job through channel j is equal to

$$\frac{\sum_i \int_{x_i^*} \tau(x, a_i) \phi_j^e(a_i, x) dx}{\sum_i \int_{x_i^*} \phi_j^e(a_i, x) dx}. \quad (18)$$

5 Quantitative Exercise

In this section, we calibrate the model to key moments in our data. This exercise generates new qualitative insights into the role of referrals in the labor market, along with new quantitative insights into how much they contribute to employment, earnings, inequality, and output. Importantly, for both sets of insights, we find that the distinction between referrals from business contacts and those from family and friends is crucial.

5.1 Parameters, Target Moments, and Identification

Maintained assumptions. We calibrate the model at a monthly frequency. Since our empirical results highlight the differential role of referrals across high- and low-skill jobs, we split our data into two sub-samples: those workers with a bachelor’s degree or more (whom we refer to as “high skill”), and those with some college or less (whom we refer to as “low skill”).²⁹ We

²⁸Note that we derive and target average wages instead of average starting wages. This is because, as is well known, the strategic wage protocol of Postel-Vinay and Robin (2002) and Cahuc et al. (2006) can often produce counterfactual starting wages for those workers hired out of unemployment. Indeed, if the option value of starting to climb the job ladder is sufficiently high, these models can even predict that workers accept negative wages early in their careers, which clearly violates constraints outside of the model (such as the minimum wage).

²⁹As we explain below, we choose to distinguish these two markets by education, as opposed to the NPB score of the occupation, to leverage existing estimates of a key parameter of this model across education groups.

think of the two markets as distinct labor markets, and hence calibrate the model separately for each skill group.

In both markets, the discount factor r is chosen to yield an annual discount rate of 95%. We also assume, in both markets, that there are two types of unobserved ability and normalize $a_1 = 1$ and $a_2 = 2$.³⁰ Finally, we choose functional forms for the production and matching technologies: we assume that the production technology is linear, $f(x) = px + c$, while the matching technologies are given by $\Gamma_j^u(x|a) = \lambda_j^u(a) [1 - H_j(x|a)]$ and $\Gamma_j^e(a) = \theta \Gamma_j^u(a)$ for $j \in \{B, F, O\}$, so that θ captures the differential arrival rates between searching on and off the job.

We assume that the rate at which workers of ability a_i meet firms through channel j satisfies $\lambda_j^u(a) = \alpha_j(a - 1) + \kappa_j$, while the distribution of match-specific productivity, $H_j(x|a_i)$, is defined by a beta distribution with shape parameters $\xi_j a$ and η . The parameter α_j captures the effect of a worker's ability on the rate at which she meets firms through channel j , while κ_j captures level differences in meeting rates across channels. The parameters ξ_j and η jointly determine the mean and variance of the distribution of match-specific productivity draws through channel $j \in \{B, F, O\}$, along with the sensitivity of these moments to workers' ability a_i .³¹

Calibration strategy. In each market (low and high skill), we set the bargaining power, β , using an estimate from outside of our sample. In particular, within the context of a model also based on Cahuc et al. (2006), Lise et al. (2016) use data from the NLSY to estimate the surplus-sharing parameter β across the same two education groups that we study. We use their estimates of $\beta = 0.188$ and $\beta = 0.272$ for the low and high skill markets, respectively.

Sixteen parameters remain for each market. To calibrate these parameter values, we calculate a vector of sixteen moments from the data, \hat{m} , and then derive the counterparts of these moments (using our analytical results) in the model, $\tilde{m}(\chi)$, for a particular vector of parameter values, χ . We then iterate over χ to minimize the loss function

$$L(\chi) = -\frac{1}{2} (\hat{m} - \tilde{m}(\chi))^T \widehat{W}^{-1} (\hat{m} - \tilde{m}(\chi)),$$

where \widehat{W} is the diagonal of the covariance matrix of \hat{m} , estimated via the nonparametric bootstrap.³² While the sixteen internally calibrated parameters are interdependent and, hence, jointly estimated, there are certain moments that are particularly informative about the value of specific parameters. Below we provide an intuitive discussion. In Appendix D we provide some more formal results regarding the identification of our model.

³⁰We also experimented with versions of the model with more than two types, but our calibration results loaded most of the weight on just two types.

³¹Given our parameterization, the mean of each distribution is $\frac{\xi_j a}{\xi_j a + \eta}$ and the variance is $\frac{\xi_j a \eta}{(\xi_j a + \eta)^2 (1 + \xi_j a + \eta)}$.

³²Appendix D provides a more detailed description of the construction of the empirical targets and the derivation of their model counterparts.

Identification. The first five parameters, which are not specific to the different job-finding channels, are most informed by relatively standard aggregate moments. The job destruction rate, δ , and the relative efficiency of on-the-job search, θ , are informed by the job-destruction (EU) and job-to-job (EE) transition rates in the data, respectively. The flow value of unemployment, b , plays a key role in determining whether workers accept job offers, and thus will be informed by the unemployment rate in each sub-sample. Lastly, aggregate wage dispersion is highly informative about the variance of match-specific productivities, which is determined in large part by the parameter η and the distribution over ability types, summarized by π_1 . Hence, to ensure that the model-generated distribution of wages matches the corresponding measure in the data, we calculate the distribution of wages in each sub-sample (high and low skill markets) after controlling for observable characteristics, and then calculate the fraction of workers in the model earning less than the wages that lie at the 25th and 75th percentiles of the empirical distribution.³³ The targets, of course, are 0.25 and 0.75, respectively.

The remaining eleven parameters must be discussed within the context of the different job-finding channels. To start, the regression results in Section 3 provide evidence consistent with lower ability workers being more likely to use F referrals and higher ability workers being more likely to use B referrals. Therefore, the average productivity of jobs found through channel F is likely to be lower, and thus the average wage of these jobs is likely to be informative about the intercept of the production technology c . Similarly, since the average productivity of jobs found through B is likely higher, the average wage of jobs found through B is informative about how output increases with the match-specific productivity, p .

The parameters $\{\kappa_B, \kappa_F, \kappa_O\}$ determine level differences in the rate at which workers contact firms through each of the three channels. As a result, κ_B and κ_F are informed by the fraction of currently employed workers who found their job through channel B and F , respectively. Since the majority of jobs are found through other channels, κ_O is most informed by the overall rate at which unemployed workers contact firms (through any channel).

The parameters $\{\alpha_B, \alpha_F, \alpha_O\}$ determine the rate at which high ability (a_2) workers contact firms through the various channels, relative to low ability (a_1) workers. Since type a_2 workers are more likely to be employed and, again, the majority of contacts occur through channel O , α_O is informed by the average contact rate of employed workers. A key mechanism in our model implies a connection between the channel that workers use to find their job and their place in the wage distribution: as we discuss in greater detail below, type a_2 workers earn higher wages and are more likely to use channel B , while type a_1 workers earn lower wages and are more likely to use channel F . Hence, following the logic in Arbex et al. (2019), we target two additional moments that capture how the fractions of employed workers who used channels B

³³Specifically, we use the same set of variables used in the regressions in Section 3: individual controls, time and region fixed effects, and the NPB score.

and F change with wages. In particular, for $j \in \{B, F\}$, we target the difference in the fraction of workers in the top wage quartile that used channel j to find their job and the fraction of workers in the bottom wage quartile that used channel j to find their job. These moments inform the parameter values of α_B and α_F .

Finally, given η , the average productivity of a match formed through channel O is determined by ξ_O , which is most informed by the average wage of workers hired through channel O . In addition to driving average wages, match-specific productivities also affect job tenure, as workers with higher productivities are less likely to leave a match. Therefore, ξ_B and ξ_F inform the average tenure of jobs found through channels B and F , respectively, relative to those found through O .

Model fit. Table 5 reports the model fit for both high and low skill markets. As one can see, the model is able to match the targeted moments quite well overall.³⁴ In Appendix F, we provide further evidence that the model is capable of replicating the empirical patterns in the data that we document in Section 3. In particular, we show that data simulated from our calibrated model generates similar relationships between the usage of referrals, the skill content of the occupation, starting wages, job tenure, and contact rates. It is important to note that, while some of these relationships are closely related to targeted moments (usage and tenure), others are not (starting wages and contact rates by channel).

5.2 Qualitative Insights from Calibration

Table 5 also summarizes the parameter values identified by the calibration. Some of these values are easy to interpret directly, and have close counterparts in related job search models. For example, the job destruction rate (δ) and the relative arrival rate of meetings for employed workers (θ) are both roughly in line with existing estimates in the literature.³⁵ The flow value of unemployment (or home production), b , is a controversial parameter in the literature, with values of this “replacement rate” ranging from 0% (Moscarini and Postel-Vinay, 2023) to 95% (Hagedorn and Manovskii, 2008) of the surplus generated from a match. Our estimates are qualitatively similar to those in Lise et al. (2016), in that we also find less educated workers

³⁴The model struggles slightly to fit the EU and EE rates in the low skill market. However, moments involving transitions between employment states are noisy in our data since we observe workers for only a year.

³⁵Cairó and Cajner (2018) provide estimates of separation rates across educational attainment levels. Our estimate of δ for workers with a college degree is slightly higher than theirs, while our estimate for workers with some college or less lies in between their (separate) estimates for workers with some college and those with only a high school degree. Similarly, the average of our estimates of θ in low and high skill markets falls within the range of recent estimates; see, e.g., Elsby and Gottfries (2022), Bilal et al. (2022), Moscarini and Postel-Vinay (2023), and Elsby et al. (2022), whose estimates of the relative intensity of on-the-job search (for all workers) lie between 0.15 and 0.2. Lise et al. (2016) arrive at estimates in a similar range as well, but interestingly find that on-the-job search is relatively more efficient for workers with less education.

Table 5: Parameter Values and Target Moments from Calibration

Parameter	High Skill	Low Skill	Closest Target	High Skill		Low Skill	
	Value	Value		Model	Data	Model	Data
δ	0.017	0.018	Job destruction (EU) rate	0.015	0.015	0.016	0.013
θ	0.248	0.121	Job transition (EE) rate	0.021	0.023	0.022	0.031
b	3.132	11.910	Unemployment rate	0.053	0.053	0.075	0.075
π_1	0.766	0.307	Fraction wages $\leq w_{25}$	0.249	0.250	0.247	0.250
η	108.536	96.141	Fraction wages $\leq w_{75}$	0.749	0.750	0.750	0.750
α_O	0.905	4.074	Contact rate employed	0.128	0.128	0.126	0.123
κ_O	0.165	0.245	Contact rate unemployed	0.258	0.244	0.348	0.350
ξ_O	9.590	41.113	Avg wage (O)	32.929	33.234	20.133	20.114
p	244.418	58.741	Avg wage (B)	36.885	36.571	21.341	21.739
c	8.263	-7.379	Avg wage (F)	32.950	32.937	20.343	20.624
α_B	1.503	1.712	Usage differential (B)	0.106	0.117	0.050	0.042
κ_B	0.051	0.036	Fraction employed (B)	0.188	0.187	0.140	0.140
α_F	0.044	0.157	Usage differential (F)	-0.026	-0.020	0.058	0.049
κ_F	0.042	0.035	Fraction employed (F)	0.204	0.204	0.273	0.274
ξ_B	8.472	40.019	Avg tenure B/O	0.853	0.853	0.887	0.887
ξ_F	11.188	50.886	Avg tenure F/O	1.158	1.158	1.298	1.305

Notes: This table reports the values of the 16 parameters that are calibrated internally, along with the values of our 16 targeted moments in the data and as computed analytically in our model, separately for the high skill and low skill markets. Most of these moments are self-explanatory. “Fraction wages $\leq w_x$ ” denotes the fraction of employed workers earning less than the wages at the $x \in \{25^{th}, 75^{th}\}$ percentiles in the data, respectively. “Usage differential (j)” refers to the difference in the fraction of workers in the top wage quartile that used channel $j \in \{B, F\}$ to find their job and the fraction of workers in the bottom wage quartile that used channel j to find their job. “Fraction employed (j)” denotes the fraction of employed workers who found their job using channel $j \in \{B, F\}$.

value leisure (or home production) more than highly educated workers.

The remaining parameters determine the properties of the channel-specific matching and production technologies—i.e., the relationships between workers’ unobserved types, the frequency with which they contact firms through different channels, and the output generated by these potential matches—that are necessary to generate the empirical patterns we observe in the data. Given the large number of these parameters, and the myriad ways they interact with one another, it is difficult to interpret these values independently. In what follows, we highlight the main qualitative properties of these technologies implied by the calibrated parameter values, explain how these properties enable the model to match the target moments in the data, and discuss the relationship between these properties and existing theories of referrals.

Key properties of the matching and production technologies. Table 6 reports a few summary statistics that reveal several key properties of the calibrated model environment. First, matching the data through the lens of our model requires significant heterogeneity in the rate at which otherwise similar workers receive offers (through any channel), both when they

are unemployed and when they are employed.³⁶ For example, in both the high and low skill markets, the overall rate at which unemployed workers of type a_2 match with firms, $\Gamma^u(x_2^*|a_2)$, is more than ten times larger than the matching rate of type a_1 unemployed workers.³⁷

The second key insight is that this heterogeneity stems mostly from differences in the arrival rate of offers through business referrals and, to a lesser extent, other (formal) channels. For example, in the high skill market, the rate at which type a_2 workers match with firms through a business referral, $\Gamma_B^u(x_2^*|a_2)$, is approximately thirty times larger than the corresponding matching rate of type a_1 workers. Offers that arrive through referrals from family and friends, in contrast, are significantly less sensitive to a worker’s unobserved type.

The third key insight is that there are important differences in expected output across channels. In both markets, F referrals create the best matches, conditional on worker type. For example, in the high skill market, the expected output for a type a_2 worker from a match formed through channel F is approximately 21% (11%) higher than the expected output from a match formed through channel B (O). This result highlights the benefits of interpreting data through the lens of a model: though one might be tempted to conclude that B referrals generate the best matches, since they are associated with higher productivity *unconditionally*, in fact F referrals generate higher productivity matches *conditional* on ability (which cannot be observed directly in the data).

Finally, while we observe similar qualitative patterns in both high and low skill markets, there are important quantitative differences. For example, while high ability workers are more likely to use B relative to F in both markets, the difference is more pronounced in the high skill market. In addition, the fraction of all meetings initiated through other channels (O) is significantly larger in the low skill market than in the high skill market. We discuss the implications of these differences in more detail below.

How do these properties generate the patterns we observe in the data? The matching rates in Table 6 are helpful for understanding how the model generates the differential usage rates of channels B and F , within and across markets. For example, in the high skill market, meetings occur relatively frequently through B , compared to the overall meeting rate, whereas meetings through B are a smaller proportion of meetings in the low skill market.

To understand how the model generates average wages and tenure conditional on job search channel, note that the difference in the arrival rate of meetings through channel B across abilities is much more pronounced than the difference in the arrival rate of meetings in channel F , i.e., $\Gamma_B(\cdot|a_2)/\Gamma_B(\cdot|a_1) > \Gamma_F(\cdot|a_2)/\Gamma_F(\cdot|a_1)$. As a result, a randomly selected employed worker who

³⁶In a related paper, Gregory et al. (2021) also find that unobserved heterogeneity in the frequency of employment transitions is an important factor in understanding labor market outcomes.

³⁷Recall that the arrival rates of offers for employed workers are simply scaled by θ , so that the statements above regarding arrival rates apply equally well to employed and unemployed workers.

Table 6: Summary Statistics from Calibrated Model

Name	Notation	a_2	High Skill			Low Skill		
			a_1	Ratio $\left(\frac{a_2}{a_1}\right)$		a_2	a_1	Ratio $\left(\frac{a_2}{a_1}\right)$
Arrival rate overall	$\lambda^u(a_i)$	2.709	0.258	10.517		6.258	0.316	19.826
Arrival rate through B	$\lambda_B^u(a_i)$	1.554	0.051	30.249		1.748	0.036	48.183
Arrival rate through F	$\lambda_F^u(a_i)$	0.085	0.042	2.053		0.191	0.035	5.511
Arrival rate through O	$\lambda_O^u(a_i)$	1.069	0.165	6.498		4.319	0.245	17.654
Expected output through B	$\int f(x)dH_B(x a_i)$	41.269	25.961	1.590		19.307	9.886	1.953
Expected output through F	$\int f(x)dH_F(x a_i)$	50.040	31.104	1.609		22.827	12.951	1.763
Expected output through O	$\int f(x)dH_O(x a_i)$	44.970	28.107	1.600		19.700	10.216	1.928

Notes: This table reports several summary statistics using the calibrated parameter values. The arrival rate is the rate of contacts with an employer. Expected output is measured as expected hourly output.

got her job through channel B is more likely to be of type a_2 , while a worker who got his job through channel F is more likely to be of type a_1 . Moreover, *ceteris paribus*, the expected output of a type a_2 worker is larger than that of a type a_1 worker and, after forming a match, type a_2 workers receive offers more frequently than type a_1 workers. Taken together, these properties imply that a randomly selected worker who got her current job through channel B will have a relatively high wage but a relatively short tenure, whereas a randomly selected worker who got his current job through channel F will have a lower wage but stay at the firm for a longer tenure.

How do these properties relate to existing theories of referrals? The empirical analysis in Section 3 revealed that distinguishing between different types of referrals is crucial for uncovering certain patterns in the data: workers who used business referrals to find their current job experienced significantly different—and, in fact, opposite—labor market outcomes, relative to workers who used a referral from a family or friend. However, the data alone do not explain how or why these two job search methods exhibit different correlations with a worker’s wage or propensity to switch jobs, or why these correlations vary systematically across occupations with different skill requirements. In this section, we exploit the results from our calibration exercise to show how interpreting the data through the lens of a structural model can help us understand the distinct roles that referrals from business contacts and family/friends play in the match formation process.

On the one hand, as we discussed above, a worker’s tendency to meet a firm through a business referral is highly sensitive to the worker’s *ex ante* type, a_i . As a result, business referrals frequently initiate contact for type a_2 workers who are likely to have high match-specific productivity at a firm, and are rarely used to initiate contact for type a_1 workers who are typically less productive. Hence, business referrals appear most consistent with theories based on a referrer’s ability to overcome *adverse selection* by screening workers based on their

expected productivity at the firm, as conjectured by Montgomery (1991) and others.

On the other hand, referrals from family and friends appear to generate relatively high productivity matches independently of a worker’s underlying type. This property is consistent with theories that ascribe a central role to the ability of referrals to generate good matches *ex post*. One such theory, put forward by Simon and Warner (1992), among others, postulates that referrals improve workers’ and firms’ ability to overcome *symmetric uncertainty* by learning about their match-specific productivity more efficiently. An alternative theory, which is also consistent with the properties of $\Gamma_F(\cdot)$ revealed by our calibration exercise, is that referrals from family and friends help overcome *moral hazard*. According to this theory (put forward by, e.g., Heath, 2018), a worker referred by a family or friend has more incentive to work hard and less incentive to shirk.

The finding that low types (a_1) rely much more heavily on referrals from friends and family is consistent with the idea that referrals can act as a “last resort” for individuals with few outside options, as espoused by Loury (2006). However, note that our findings are *not* consistent with theories of referrals based on nepotism or other forms of favoritism, in which referred workers are, on average, less productive than workers hired through more formal channels (see, e.g., Bandiera et al., 2009; Fafchamps and Moradi, 2015).

5.3 Quantitative Insights from Calibration

In this section, we use the calibrated model to explore the quantitative effects of referrals from business contacts and family and friends on employment, earnings, output, and inequality. To do so, we simulate the labor market experience of a cohort of workers who enter the market unemployed at $t = 0$, assuming they have access to all three job search channels, and follow them for a period of ten years.³⁸ Then, we repeat the simulation but shut down referrals from business contacts, by setting $\Gamma_B(x|a) = 0$, and evaluate the labor market outcomes of different types of workers (a_1 and a_2) across low- and high-skill labor markets. Finally, we repeat the exercise but shut down referrals from family and friends, while restoring a worker’s ability to meet firms through business contacts.³⁹

The contribution of referrals to earnings. Table 7 reports the average annual earnings, the average employment rate, and the average wage (conditional on being employed) of workers

³⁸The results reported below are similar at 20 and 30 year horizons, as well.

³⁹Note that, in these quantitative exercises, we assume that shutting down one channel has no effect on the arrival rate or quality of matches generated by other channels. This allows for a clean decomposition of the relative contribution of each channel to labor market outcomes. Moreover, existing evidence shows that workers already exploit available job search channels, with plenty of time to spare; see, e.g., Mukoyama et al. (2018). This suggests that workers without access to one particular job search channel may not be able to easily increase the arrival rate and/or quality of matches through an alternative channel.

Table 7: Contribution of Referrals to Earnings and Employment

Outcome	Exercise	High Skill			Low Skill		
		Low Ability	High Ability	Overall	Low Ability	High Ability	Overall
Annual Earnings	Benchmark	\$45,196	\$100,460	\$58,244	\$18,348	\$45,650	\$37,207
	Shut down B Referrals	-5.4%	-5.0%	-5.3%	-2.0%	-1.1%	-1.0%
	Shut down F Referrals	-7.4%	-0.8%	-4.9%	-11.2%	-1.2%	-2.6%
Employment Rate	Benchmark	0.89	0.97	0.91	0.67	0.96	0.87
	Shut down B Referrals	-2.1%	-0.9%	-1.8%	-2.0%	-0.2%	-0.5%
	Shut down F Referrals	-1.4%	0.0%	-1.1%	-9.2%	0.2%	-1.9%
Wage	Benchmark	\$24.21	\$49.79	\$30.25	\$13.17	\$22.74	\$19.79
	Shut down B Referrals	-3.5%	-4.1%	-3.8%	0.1%	-0.9%	-0.5%
	Shut down F Referrals	-6.2%	-0.8%	-4.3%	-2.2%	-1.4%	-1.4%

Notes: This table reports estimates from a simulation of our calibrated model in which a cohort of 30,000 workers enters the market unemployed. We follow this cohort for a period of 10 years. The top panel reports average annual earnings over these 10 years in both markets (high and low skill) for low and high ability workers (a_1 and a_2 , respectively). The middle panel reports the average employment rates for each of these groups, and the bottom panel reports average hourly wages of employed workers. In each panel, the top row reports results from our benchmark model, while the second and third rows report how the baseline results change when we shut down referrals from business contacts and family/friends, respectively.

during their first ten years in the labor market.⁴⁰ Several interesting insights emerge.

First, and perhaps most striking, is the extent to which low ability workers depend on referrals from family and friends. For instance, our simulations suggest that more than 11% (7%) of earnings of low ability workers in low (high) skill occupations can be attributed to referrals from family and friends. Intuitively, referrals from family and friends are critically important to this subset of workers for two reasons: first, because they struggle to contact firms through other channels (O or B); and second, because matches formed through family and friends tend to be high quality matches (for all abilities). Second, while our calibration suggests that low ability workers depend heavily on referrals from family and friends, high ability workers do not: in both low and high skill markets, high ability workers match with firms through channels B and O sufficiently frequently that shutting down channel F has very modest effects on earnings and, in fact, no effect on employment.

The effects of referrals from business contacts are quite different: while the contribution of channel B to labor market outcomes is more even across workers' abilities, our calibration suggests that workers' networks of business contacts are considerably more important in high skill occupations than in low skill occupations. For example, shutting down channel B reduces earnings by about 5% in the former, and only about 1% in the latter. Intuitively, there are several reasons why the effects of business referrals are more pronounced in high skill occupations. For one, referrals from business contacts are used more frequently in the high skill market, especially among high ability workers who have higher average match-specific produc-

⁴⁰Note that the average employment rate reported in Table 7 is lower than in the steady state because, in the simulation, each individual starts off unemployed.

tivity. Moreover, since the calibrated production function is steeper in the high skill market, a reduction in average match quality translates into a larger loss in output.⁴¹

Finally, since the quantitative effect of referrals differs significantly both across B and F , and across low and high skill occupations, Table 7 illustrates again the importance of distinguishing between different types of referrals and the different types of occupations in which they are used. Importantly, these results also highlight the advantage of interpreting the data through the lens of a model. For instance, the regression results in Table 2, which report the correlation between the use of these two types of referrals and wages, might lead one to believe that referrals from business networks have a large, positive effect on workers' wages, while referrals from family and friends have modest, negative effects on wages. However, by allowing for unobserved heterogeneity and accounting for selection effects, our modeling exercise illustrates that this conclusion would be erroneous. Instead, our calibration reveals that referrals from family and friends have a large, positive effect on workers' wages, but they tend to be used by workers who otherwise struggle to find jobs (and hence earn, on average, lower wages).

Referrals, output, and inequality. While referrals are widely regarded as an important channel for matching workers and firms, a central, yet open question is whether referrals mitigate or exacerbate earnings inequality. More specifically, do referrals primarily help workers with dim employment prospects to find good-paying jobs? Or, instead, do they mostly help relatively well-connected, high-wage workers to earn even more? The answers to these questions are key for assessing the implications of referrals on earnings inequality and evaluating an explicit role for referrals in the context of labor market policies.

Table 8 reports the change in total output and inequality (as measured by the standard deviation of earnings) in the simulations described above, in which we separately shut down B and F in low and high skill labor markets. The table reveals that the relationship between referrals and inequality again relies on the crucial distinction between referrals from different sources. In particular, the simulations reveal that referrals from business contacts increase earnings inequality, particularly in high skill occupations, by increasing the wages of high ability workers, who are already well-paid. Referrals from family and friends, however, do the exact opposite: they reduce earnings inequality, particularly in low skill occupations, by increasing employment and wages among low ability workers (while doing very little for high ability,

⁴¹In general, shutting down any channel $j \in \{B, F, O\}$ has three effects: workers spend more time in unemployment; they have lower match-specific productivity, on average, when matched; and they receive a smaller share of the surplus, since outside offers arrive more slowly. In most cases, the reduction in earnings associated with shutting down job search channels is driven predominantly by a reduction in the wages of employed workers—both because they have lower match-specific productivity and because they get fewer outside offers. However, in the case of low ability workers in the low skill market, it is the opposite—the effect is mainly caused by lower levels of employment. This is driven by the fact that low ability workers, in particular in the low skill market, have a much harder time meeting firms than do high skill workers.

high-wage workers).

Hence, the central question of whether referrals generate a trade-off between output and inequality depends crucially on the distinction between these two types of referrals. In general, employers almost always allow—and, indeed, encourage—their employees to refer candidates from their business or social network. However, out of concerns for nepotism, some employers do *not* allow current employees to refer their relatives. Interestingly, our results suggest that these types of restrictions may actually generate less output and greater inequality, as they do the most damage to workers who struggle to find work through other channels.

Table 8: Contribution of Referrals to Output and Inequality

Exercise	Total Output		St. Dev. of Earnings	
	High Skill	Low Skill	High Skill	Low Skill
Benchmark	\$70,695	\$38,469	\$25,319	\$13,103
Shut down B Referrals	-2.9%	-0.7%	-2.7%	-0.5%
Shut down F Referrals	-3.8%	-2.8%	3.8%	5.3%

Notes: This table reports estimates from a simulation of our calibrated model in which a cohort of 30,000 workers enters the market unemployed. We follow this cohort for a period of 10 years. The left panel measures average (hourly) output over these 10 years, and the right panel measures the standard deviation in earnings across workers. In the first row we report these measures separately by market and by ability type. In the second and third rows we report how these baseline results change when we shut down referrals from business contacts and family/friends, respectively.

6 Concluding Remarks

A longstanding challenge in labor economics is understanding the channels through which workers and firms form productive matches, and the implications for labor market outcomes. In surveys of both workers and firms, referrals are often cited as a key input into the match formation process. However, while there may be a consensus that referrals are widely used, there is far less agreement about *who* uses them most frequently, *why* they are valuable, or *how* they ultimately affect match quality, wages, turnover, inequality, and output.

In this paper, we try to make progress on these important questions. Our contribution can be broken down into three parts. First, leveraging a relatively new dataset, we show that clear patterns emerge from the data only after distinguishing between different types of referrals and different types of jobs. Second, by interpreting these patterns through the lens of a model, we are able to assess various theories of referrals and explore which mechanisms are (or are not) consistent with the patterns we find in the data. Lastly, by further exploiting our calibration results, we are able to derive quantitative estimates of the contribution of referrals from different

sources in low- and high-skill labor markets.

We find that referrals from friends and family are a crucial source of jobs for a subset of workers that struggle to generate offers and matches through more traditional channels. Indeed, this type of referral improves earnings and employment outcomes, and represents an important force for reducing earnings inequality. In contrast, business referrals tend to increase the wages of workers who have relatively good employment prospects to begin with. Hence, the use of business referrals, which is typically encouraged by firms, increases output but also exacerbates earnings inequality. These findings are important in assessing the impact of referrals in the labor market and in evaluating their role within the broader context of labor market policies.

While this paper focuses on the direct effects of referrals on labor market outcomes, we believe our results could deepen our understanding of other economic phenomena as well. As a leading example, consider the literature that studies worker mobility across geographic locations. Within this literature, one key finding is that substantial moving costs are needed to rationalize the internal migration patterns observed in the data, particularly for low-skill workers (Kennan and Walker, 2011; Diamond, 2016; Piyapromdee, 2021).⁴² In addition, the literature studying the Moving to Opportunity programs has found that, despite previous research documenting the importance of neighborhoods for economic outcomes, moving families from high-poverty to low-poverty neighborhoods has no significant effect on earnings or employment (Katz et al., 2001; Chetty et al., 2016). Since we show that low-skill workers rely heavily on referrals from friends and relatives, and relocating often severs a worker’s connections to many contacts in his social network, our results provide a natural explanation for—and a bridge between—these two important findings. Thus, by offering a deeper explanation of the *sources* of migration costs, our findings have potentially important implications for the design of social programs aimed at improving labor market outcomes by promoting migration. Analogously, our results on referrals from business contacts could provide new insights into the costs associated with occupational mobility, as changing occupations may significantly alter the structure of one’s business network. We leave exploring and deepening these connections for future work.

⁴²Indeed, Zerecero (2021) documents particularly high migration costs away from one’s birthplace.

A Data Description

In this Appendix we describe how we arrive at our final estimation sample and provide additional details regarding the construction of our wage and tenure variables. We also provide a list of 2-digit occupation codes and their corresponding NPB scores.

Our data set combines the annual Job Search supplement of the SCE from 2013-2018. We keep individuals that are of working age (18 to 64) and that are not self-employed, for a total of 5,099 observations. After excluding individuals who work in military occupations and dropping a small number of observations with missing demographic data or inconsistent wage data, we are left with a final sample of 5,062 observations.

All three wage measures (current, starting, previous) are reported as either hourly, weekly, or annual. Survey respondents are also asked to report their usual hours spent at their job per week for both their current job and their previous job.⁴³ We divide weekly wages by usual hours and annual wages by usual hours and by 52 to convert everything to hourly wages. We then deflate all three nominal wage measures using the relevant CPI index obtained from the BLS.

For job tenure, the SCE survey asks workers the month and year in which they started their current job. We use this information to compute the duration of the worker's current job at the time of the survey.

⁴³For the 2013 data, survey respondents were not directly asked the usual hours on their previous job, and instead were asked how much their hours increased or decreased from their previous job. We use this change and the reported usual hours at the current job to construct previous hours.

Table 9: Nam-Powers-Boyd Index (2-Digit Occupation Level)

Occupation	NPB Index
Food Preparation and Serving Related Occupations (FOOD)	17
Building and Grounds Cleaning and Maintenance Occupations (BLDG)	17
Personal Care and Service Occupations (PERS)	27
Transportation and Material Moving Occupations (TRSP)	32
Production Occupations (PROD)	33
Construction and Extraction Occupations (CSTR)	34
Healthcare Support Occupations (NURS)	39
Sales and Related Occupations (SLS)	43
Office and Administrative Support Occupations (ADMN)	47
Installation, Maintenance, and Repair Occupations (MNT)	47
Protective Service Occupations (PROT)	55
Arts, Design, Entertainment, Sports, and Media Occupations (ART)	64
Community and Social Service Occupations (SOC)	72
Education, Training, and Library Occupations (EDU)	75
Healthcare Practitioners and Technical Occupations (DOC)	78
Business and Financial Operations Occupations (BUS)	81
Life, Physical, and Social Science Occupations (LIFE)	83
Management Occupations (MGT)	84
Architecture and Engineering Occupations (ENG)	86
Computer and Mathematical Occupations (COMP)	87
Legal Occupations (LEGL)	88

Notes: This table provides the Nam-Powers-Boyd (NPB) occupational index score aggregated to the 2-digit occupation level.

B Additional Empirical Results

Job Satisfaction and Referrals

Table 10 reports results from linear regressions relating the job search method used to find a worker’s current job and their reported satisfaction with that job in response to the following questions:

1. “Taking everything into consideration, how satisfied would you say you are, overall, in your current job?”
2. “How satisfied would you say you are with your level of compensation at your current job?”
3. “How satisfied would you say you are with other aspects of the job, such as benefits, maternity/paternity leaves, flexibility in work hours, etc?”
4. “How well do you think this job fits your experience and skills?”
5. “How would you rate the opportunities for a promotion or other career progression with your current employer, over the next three years?”

For the first three measures respondents are asked to respond on a scale from 1-5 capturing “very dissatisfied” to “very satisfied”. For the last two, they are asked to respond on a scale from 1-7 ranging from “very poor” to “very good”. As the results indicate, there are no systematic differences in job satisfaction across job search method.

Search Behavior of Employed Workers and Referrals

Table 11 reports results from regressions relating four measures of on-the-job search to the method used to find a worker’s current job. The measures we use are number of job applications sent out in the past 4 weeks, whether any job search was performed over the past weeks as well as the past 12 months, and the number of hours of job search over the past 4 weeks. Overall there do not seem to be any large differences in on-the-job search based across job finding method. There is some weak evidence that individuals who found their job via family/friends referral are slightly less likely to have searched over the past 12 months, although not over the past 4 weeks. There is also some weak evidence that business referred workers search more in terms of hours, but not in terms of overall probability.

Table 10: Job Satisfaction and Referrals

	Job Satisfaction				
	Overall	Compensation	Other	Fit	Promotion
Business Referral	0.080 (0.051)	0.069 (0.054)	0.093* (0.053)	0.069 (0.059)	0.132 (0.089)
Family/Friends Referral	0.019 (0.045)	0.013 (0.048)	-0.032 (0.048)	0.013 (0.052)	-0.050 (0.080)
Skill Index	0.004*** (0.001)	0.005*** (0.001)	0.006*** (0.001)	0.007*** (0.001)	0.005*** (0.002)
Time and Region FE	✓	✓	✓	✓	✓
Individual Controls	✓	✓	✓	✓	✓
N	3068	3067	3067	3067	3068

Notes: Estimates are from regressions regarding five different measures of job satisfaction. Individual controls include age, gender, race, marital status, number of children under the age of 6, and home ownership status. Standard errors are presented in parentheses below the point estimates. *** denotes significance at the 1% level, ** denotes significance at the 5% level, and * denotes significance at the 10% level. These questions were not asked in 2013, which reduces the number of observations by 711. For three of the measures, there is one individual for whom we do not observe answers regarding job satisfaction.

C Omitted Proofs

Proof of Proposition 1

Using the thresholds described in the text, we can write the flow Bellman equation characterizing the value of unemployment for a worker with ability a_i as

$$rV^u(a_i) = b + \sum_{j \in \mathcal{C}} \lambda_j^u(a_i) \int_{x_i^*}^{\bar{x}} [V^e(a_i, x, w^u(a_i, x)) - V^u(a_i)] dH_j(x|a_i).$$

Since $d\Gamma^u(x|a_i) = -\sum_{j \in \mathcal{C}} \lambda_j^u(a_i) dH_j(x|a_i) < 0$, this expression simplifies to

$$[r + \Gamma^u(x_i^*|a_i)] V^u(a_i) = b - \int_{x_i^*}^{\bar{x}} V^e(a_i, x, w^u(a_i, x)) d\Gamma^u(x|a_i). \quad (19)$$

In words, as in standard job search models, the value of unemployment is equal to the flow value of unemployment, b , plus the option value of finding a job out of unemployment.

Table 11: On-the-Job Search and Referrals

	On-the-Job Search			
	# of Applications (Last 4 Weeks)	Any Search (Last 4 Weeks)	Any Search (Last 12 Months)	Search Hours (Last 4 Weeks)
Business Referral	-0.028 (0.217)	0.020 (0.018)	0.005 (0.022)	0.252* (0.149)
Family/Friends Referral	-0.190 (0.195)	0.010 (0.016)	-0.036* (0.020)	0.148 (0.134)
Skill Index	-0.012*** (0.004)	-0.001** (0.000)	-0.000 (0.000)	-0.010*** (0.003)
Time and Region FE	✓	✓	✓	✓
Individual Controls	✓	✓	✓	✓
N	3779	3765	3118	3779

Notes: Estimates are from regressions on the total number of applications sent in the past 4 weeks, indicators for whether individuals have searched at all for a job over the past 4 weeks and past 12 months, and the total number of job search hours over the past 4 weeks. Individual controls include age, gender, race, marital status, number of children under the age of 6, and home ownership status. Standard errors are presented in parentheses below the point estimates. *** denotes significance at the 1% level, ** denotes significance at the 5% level, and * denotes significance at the 10% level. There are 14 observations for which we do not observe search behavior over the past 4 weeks. For search behavior over the past 12 months we exclude observations for which individuals have a tenure of less than 12 months, since we cannot determine whether they were searching on or off the job.

Following similar steps for employed workers reveals that

$$\begin{aligned}
[r + \delta + \Gamma^e(\hat{x}|a_i)] V^e(a_i, x, w) = & w + \delta V^u(a_i) - \int_{\hat{x}}^x V^e(a_i, x, w^e(a_i, x', x)) d\Gamma^e(x'|a_i) \\
& - \int_x^{\bar{x}} V^e(a_i, x', w^e(a_i, x, x')) d\Gamma^e(x'|a_i), \tag{20}
\end{aligned}$$

where $\hat{x} \equiv \hat{x}(a_i, x, w)$.

By construction, the $\hat{x}(a, x, f(x))$ threshold has the following properties, which are useful later:

$$x = \hat{x}(a, x, f(x)) \tag{21}$$

$$V^e(a_i, x, w) = V^e(a_i, \hat{x}(a_i, x, w), f(\hat{x}(a_i, x, w))). \tag{22}$$

Now, by substituting (8), equation (20) can be written

$$\begin{aligned}
[r + \delta + \Gamma^e(\hat{x})] V^e(a_i, x, w) = & w + \delta V^u(a_i) + V^e(a_i, x, f(x)) [\beta [\Gamma^e(\hat{x}|a_i) - \Gamma^e(x|a_i)] + (1 - \beta) \Gamma^e(x|a_i)] \\
& - (1 - \beta) \int_{\hat{x}}^x V^e(a_i, x', f(x')) d\Gamma^e(x'|a_i) - \beta \int_x^{\bar{x}} V^e(a_i, x', f(x')) d\Gamma^e(x'|a_i),
\end{aligned}$$

where, again, we have used $\hat{x} \equiv \hat{x}(a_i, x, w)$ to economize on notation.

Setting the wage $w = f(x)$, substituting (21), and simplifying yields

$$[r + \delta + \beta \Gamma^e(x)] V^e(a_i, x, f(x)) = f(x) + \delta V^u(a_i) - \beta \int_x^{\bar{x}} V^e(a_i, x', f(x')) d\Gamma^e(x'|a_i).$$

Differentiating with respect to x then yields

$$\frac{\partial V^e(a_i, x, f(x))}{\partial x} = \frac{p}{r + \delta + \beta \Gamma^e(x|a_i)}. \quad (23)$$

Using this relationship in the expression for $V^e(a_i, x, w)$ above, integrating by parts, using (20), and simplifying yields

$$(r + \delta) V^e(a_i, x, w) = w + \delta V^u(a_i) + p \left[(1 - \beta) \int_{\hat{x}}^x \frac{\Gamma^e(x'|a_i)}{r + \delta + \beta \Gamma^e(x'|a_i)} dx' + \beta \int_x^{\bar{x}} \frac{\Gamma^e(x'|a_i)}{r + \delta + \beta \Gamma^e(x'|a_i)} dx' \right]. \quad (24)$$

Plugging in $w = w^e(a_i, x, x')$, subtracting $(r + \delta) V^e(a_i, x, f(x))$, and using (20) yields

$$(r + \delta) [V^e(a_i, x', f(x')) - V^e(a_i, x, f(x))] = w^e(a_i, x, x') - f(x) + p \left[(1 - \beta) \int_x^{x'} \frac{\Gamma^e(x''|a_i)}{r + \delta + \beta \Gamma^e(x''|a_i)} dx'' - \beta \int_x^{x'} \frac{\Gamma^e(x''|a_i)}{r + \delta + \beta \Gamma^e(x''|a_i)} dx'' \right]$$

since $\hat{x}(a_i, x, w^e(a_i, x, x')) = x$. Using (23) yields

$$w^e(a_i, x, x') = (r + \delta) \beta \int_x^{x'} \frac{p}{r + \delta + \beta \Gamma^e(x''|a_i)} dx'' + f(x) - p(1 - 2\beta) \int_x^{x'} \frac{\Gamma^e(x''|a_i)}{r + \delta + \beta \Gamma^e(x''|a_i)} dx''.$$

Straightforward algebra then yields the expression in (8).

Next, to characterize x_i^* , note that we can use (2) and (23) to get that

$$\begin{aligned} r V^u(a_i) &= b - \int_{x_i^*}^{\bar{x}} [V^e(a_i, x, w^u(a_i, x)) - V^u(a_i)] d\Gamma^u(x|a_i) \\ &= b - \beta \int_{x_i^*}^{\bar{x}} [V^e(a_i, x, f(x)) - V^u(a_i)] d\Gamma^u(x|a_i) \\ &= b + \beta \int_{x_i^*}^{\bar{x}} \frac{p \Gamma^u(x|a_i)}{r + \delta + \Gamma^e(x|a_i)} dx, \end{aligned} \quad (25)$$

where the last equality follows from integration by parts. We can also use (1) and (24) to write

$$\begin{aligned} rV^e(a_i, x_i^*, f(x_i^*)) &= f(x_i^*) + \delta V^u(a_i) + \beta \int_{x_i^*}^{\bar{x}} \frac{p\Gamma^e(x|a_i)}{r + \delta + \Gamma^e(x|a_i)} dx \\ \Rightarrow rV^u(a_i) &= f(x_i^*) + \beta \int_{x_i^*}^{\bar{x}} \frac{p\Gamma^e(x|a_i)}{r + \delta + \Gamma^e(x|a_i)} dx. \end{aligned} \quad (26)$$

Equating (25) and (26) yields (10).

Finally, to characterize the equilibrium distributions, note that substituting (5) into (6) yields (11), and substituting (11) into (7) then yields (12).

Proof of Lemma 1

Let $G(w|a_i, x)$ denote the fraction of type a_i workers currently employed at a firm with productivity x that earn a wage $w' \leq w$. In a steady-state equilibrium, the outflow of such workers is

$$G(w|a_i, x)\phi^e(a_i, x) [\delta + \Gamma^e(\hat{x}(a_i, x, w)|a_i)]. \quad (27)$$

Intuitively, the product of the first two terms yields the measure of workers of type a_i employed at a firm with productivity x that earn a wage $w' \leq w$. The third term yields the rate at which these workers exit the set, either because the job is destroyed or because they contact a new firm with productivity $x' > \hat{x}(a_i, x, w)$.

The inflow of workers into this set can be written

$$\phi^u(a_i) \sum_j \lambda_j^u(a_i) h_j(x|a_i) + \Phi^e(\hat{x}(a_i, x, w)|a_i) \sum_j \lambda_j^e(a_i) h_j(x|a_i). \quad (28)$$

Intuitively, type a_i individuals in employment state $k \in \{e, u\}$ receive opportunities to be employed at firms of type x through channel j at rate

$$\sum_j \lambda_j^k(a_i) h_j(x|a_i) = -d\Gamma^k(x|a_i).$$

However, they will only accept and earn a wage $w' \leq w$ if (i) they are hired from unemployment, or (ii) they were employed at a firm of type $x' < x$ such that $w^e(a_i, x', x) \leq w$ or, equivalently, a firm of type $x' \leq \hat{x}(a_i, x, w)$. Equating the outflow and inflow in equations (27) and (28) yields the result.

For the sake of completeness, here we derive $g(w|a_i, x) = dG(w|a_i, x)$, the density of wages

across workers of ability a_i currently employed at productivity x . Differentiating (13) yields

$$g(w|a_i, x) = -\frac{\partial \hat{x}}{\partial w} \left\{ \frac{\phi^e(\hat{x}|a_i) d\Gamma^e(x|a_i) [\delta + \Gamma^e(\hat{x})] - [\phi^u(a_i) d\Gamma^u(x|a_i) + \Phi^e(\hat{x}|a_i) d\Gamma^e(x|a_i)] d\Gamma^e(\hat{x})}{\phi^e(a_i, x) [\delta + \Gamma^e(\hat{x})]^2} \right\},$$

where

$$\frac{\partial \hat{x}}{\partial w} = \frac{r + \delta + \beta \Gamma^e(\hat{x}|a_i)}{p(1 - \beta) [r + \delta + \Gamma^e(\hat{x}|a_i)]}.$$

D Quantitative Exercise

Moments used in the calibration

Here we provide a more detailed description of the sixteen moments we use to calibrate the model.

Unemployment rate. The unemployment rate implied by the model is $u = \sum_i \phi^u(a_i)$, where $\phi^u(a_i)$ is characterized in equation (11).

EU rate. The job destruction rate in the continuous time model is simply $1 - e^{-\delta}$. However, given the relatively high job-finding rate, Shimer (2005) argues that it's important to account for time aggregation. Hence, we adjust the separation rate to account for time aggregation as Shimer (2005) suggests in equation (2) (on page 32), so that

$$\text{separation rate} = (1 - e^{-\delta}) (1 - 0.5 \times [\text{job finding rate}])$$

where the job finding rate is determined endogenously:

$$\sum_i \frac{\Phi^u(a_i)}{u} [1 - e^{-\Gamma^u(x_i^*|a_i)}].$$

EE rate. A worker currently matched at a type x job moves if and only if he contacts a firm and draws $x' > x$. Hence, the job-to-job transition rate is:

$$\sum_i \underbrace{\frac{\Phi^e(1|a_i)}{1-u}}_{\text{frac type } i} \underbrace{\int_{x_i^*}^{\bar{x}} [1 - e^{-\Gamma^e(x|a_i)}] \frac{d\Phi^e(x|a_i)}{\Phi^e(1|a_i)}}_{\text{prob type } i \text{ moves}} = \frac{1}{1-u} \sum_i \int_{x_i^*}^{\bar{x}} [1 - e^{-\Gamma^e(x|a_i)}] d\Phi^e(x|a_i).$$

Contact rates. The model is silent on whether a meeting between a worker and firm that has a negative surplus (i.e., when $x < x_i^*$) should be counted as a “contact.” We adopt the convention that unemployed workers count every meeting as a contact, and hence the average rate at which unemployed workers contact a firm is

$$\frac{\sum_i \phi^u(a_i)}{u} [1 - e^{-\Gamma^u(0|a_i)}].$$

Alternatively, we assume that an employed worker is uninterested in initiating a contact with a new firm if the surplus is negative. Hence, the arrival rate of contacts for employed workers is

$$\sum_i \frac{\Phi^e(1|a_i)}{1-u} [1 - e^{-\Gamma^e(x_i^*|a_i)}].$$

Residual wage dispersion. Since $G(w|a_i, x)$ denotes the fraction of type a_i workers who are employed at a firm with match-specific productivity x earning a wage $w' \leq w$, the fraction of *all* workers earning $w' \leq w$ is

$$\frac{\sum_i \int_{x_i^*}^{\bar{x}} G(w|a_i, x) \phi^e(a_i, x) dx}{1-u}.$$

Plugging in the closed form expression for $G(w|a_i, x)$ in (13) allows us to calculate the fraction of employed workers earning less than the wages that lie at the 25th and 75th percentiles of the empirical distribution.

Usage across channels. The fraction of employed workers who found their job through channel $j \in \{B, F\}$ is characterized in equation (15).

Average wages across channels. The average wage of employed workers who found their job through channel $j \in \{B, F, O\}$ is characterized in equations (16) and (17).

Average tenure across channels. The average tenure of employed workers who found their job through channel $j \in \{B, F, O\}$ is characterized in equation (18).

Differential usage of channels across the wage distribution. Let $\omega_j(a, x|w)$ denote the cumulative measure of workers of type a who are matched at a firm with productivity x that got their job through channel $j \in \{B, F, O\}$ and currently earn wage $w' \leq w$. In a stationary equilibrium we must have

$$\begin{aligned} \dot{\omega}_j(a, x|w) &= \Phi^e(\hat{x}(a, x, w)|a) d\Gamma_j^e(x|a) + \phi^u(a) d\Gamma_j^u(x|a) \mathbf{1}_{\{w \geq w^u(x, a)\}} \\ &\quad - \omega_j(a, x|w) [\delta + \Gamma^e(\hat{x}(a, x, w)|a)] = 0. \end{aligned}$$

The first line in the expression above captures the inflow of workers into the set. First, a mass $\Phi^e(\hat{x}(a, x, w)|a) d\Gamma_j^e(x|a)$ of type a workers accept offers that arrived through channel j at a job with match-specific productivity x and earn a wage less than or equal to w . Second, a mass $\phi^u(a) d\Gamma_j^u(x|a)$ of type a unemployed workers accept an offer that arrived through channel j at a job with match-specific productivity x , and hence enter the set if $w \geq w^u(a, x)$. The second

line in the expression captures the outflow of workers, who exit either because the match is destroyed or because an offer arrives that increases their current wage above w (either at the incumbent or the poaching firm). Solving yields

$$\omega_j(a, x|w) = \frac{\Phi^e(\hat{x}(a, x, w)|a) d\Gamma_j^e(x|a) + \phi^u(a) d\Gamma_j^u(x|a) \mathbf{1}_{\{w \geq w^u(x, a)\}}}{\delta + \Gamma^e(\hat{x}(a, x, w)|a)}.$$

Using this, the measure of workers earning wage $w' \leq w$ that got their job through channel j is:

$$\Omega_j(w) = \sum_i \int \omega_j(a_i, x|w) dx.$$

Then the fraction of workers earning $w' \leq w$ that got their job through channel j is:

$$\frac{\Omega_j(w)}{\sum_j \Omega_j(w)}. \quad (29)$$

To construct the final two moments, we use (29) to calculate the fraction of workers in the top and bottom quartiles of the wage distribution that found their job through channel $j \in \{B, F\}$, and then we take the difference.

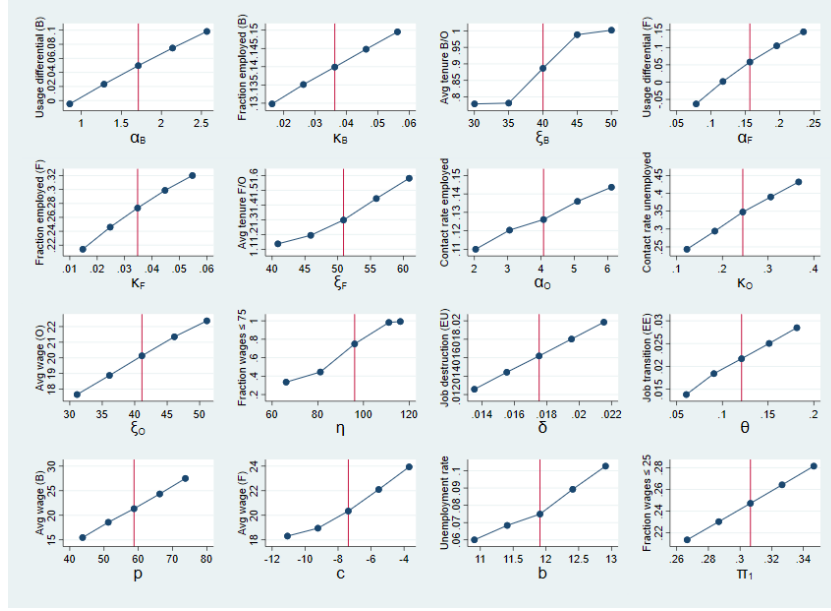
E Identification

In this Appendix we report the results of two sets of identification exercises designed to demonstrate the identification of the model's parameters. We report results corresponding to the low skill market, but we find similar results for the high skill market, and those results are available upon request.

As we discuss in the main text, we match each moment to a (unique) corresponding parameter that is particularly relevant for that moment. In our first exercise, for each moment, we vary the assigned parameter around the optimal value, while holding the other parameters fixed at their optimal values, and compute how the moment varies. Figure 2 presents the results of this exercise. As the figure demonstrates, all of the moments demonstrate significant monotonic variation with respect to their corresponding parameters.

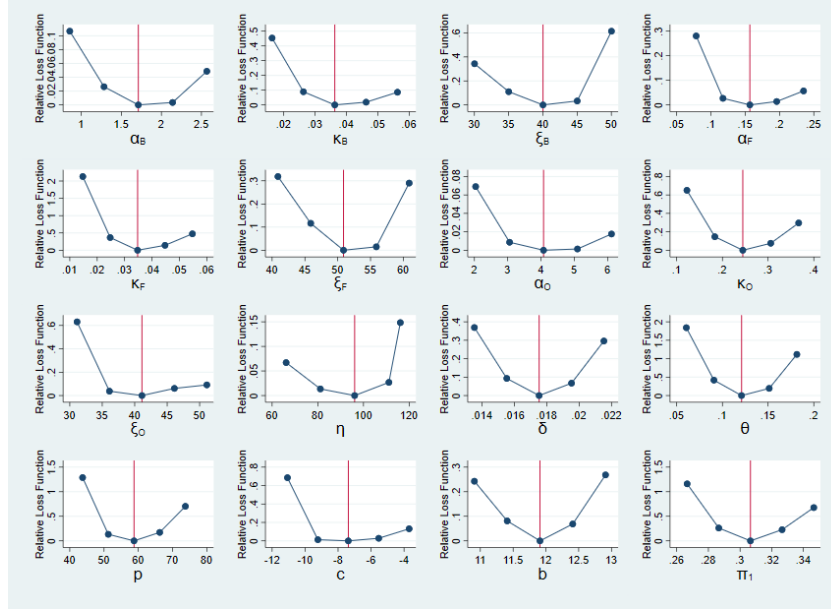
The second exercise we perform is to vary each parameter around the calibrated value, re-calibrate all the other parameters, and then report the relative difference in the loss functions. The results of this exercise are reported in Figure 3. We find that the loss functions are u-shaped around the optimal values of each parameter, providing additional evidence that our model is well identified.

Figure 2: Moments as a Function of Assigned Parameter



Notes: For each of the parameters, we plot how each moment varies with its assigned parameter, while holding all other parameters fixed at their optimal values. The calibrated value for each parameter is indicated by the red line.

Figure 3: Relative Loss Function



Notes: This figure plots how the loss function changes when we fix each parameter at values around the calibrated value and re-calibrate all of the other parameters. The vertical axis plots the relative difference in the loss functions. The optimal value of the parameter is indicated by the red line.

F Regressions on Simulated Data

In this Appendix we show that our model is capable of replicating the empirical patterns we demonstrate in the SCE data in Section 3 which relate referrals usage to various labor market outcomes. In particular, we replicate the regression results in that section (see Tables 1 - 4) on our simulated data. For convenience we also report the corresponding results from Section 3.

We first examine the relationship between the usage of referrals and the skill content of the occupation. In the real data we use a continuous measure of skill, the NPB score. Since in the simulated data our measure of occupational skill is binary, we replicate the results in Table 1 using the binary measure for the skill content of the occupation as defined in Section 5. Table 12 shows that our model simulated data generates a very similar relationship as compared to the real data for both business and family/friends referrals.

Table 12: Simulated Regressions Usage

	<i>Model</i>			<i>Data</i>	
	Type of Referral			Type of Referral	
	Business	Family/Friends		Business	Family/Friends
Skill / Market	0.051*** (0.012)	-0.075*** (0.014)		0.043*** (0.013)	-0.076*** (0.015)

Notes: Estimates are from regressions in which the outcome is whether an individual used either a business or family/friend referral to find their current job. Results in the first set of columns correspond to estimates based on data simulated from our model. Results in the second set of columns correspond to estimates based on the SCE data. For the simulated data regressions, we draw a sample size equal to that used in the real data. We repeat this 100 times and report the means and standard deviations of those estimates. Standard errors are presented in parentheses below the point estimates. *** denotes significance at the 1% level, ** denotes significance at the 5% level, and * denotes significance at the 10% level.

Next we examine how starting wages, job tenure, and contact rates vary with the channel used to find the job. For each set of results we report regressions run on data simulated from our model as well as the SCE data. For the results using the simulated data we use our binary measure of skill. For the results using the SCE data, we use the continuous measure of skill, and thus these results are identical to those reported in Section 3, but we reproduce them again here for convenience.⁴⁴

In Table 13 we look at the relationship between starting wages and the usage of both types of referrals. For business referrals, we find a positive and statistically significant association with starting wages that decreases in magnitude once we control for previous wages, a rough proxy

⁴⁴We have also run these regressions in the real data using the binary measure of skill, and the estimated relationships with referrals usage are quantitatively very similar.

for the unobserved ability of the worker, as we find in the real data. For family/friends referrals, we find a negative relationship when not controlling for the skill content of the occupation or the previous wage, as in the real data, although the magnitude is smaller, and as a result not statistically significant. We also find that controlling for previous wage causes this estimate to increase, leading to a small and statistically insignificant relationship between family/friends referrals, as in the real data. The simulated regression results also deliver a very similar positive correlation between previous wages and starting wages.

Table 13: Simulated Regressions Starting Wages

	<i>Model</i>				<i>Data</i>		
	Log Real Starting Wage				Log Real Starting Wage		
	(1)	(2)	(3)		(1)	(2)	(3)
Business Referral	0.122*** (0.024)	0.107*** (0.021)	0.033** (0.017)		0.161*** (0.025)	0.148*** (0.023)	0.085*** (0.022)
Family/Friends Referral	-0.024 (0.016)	0.000 (0.016)	0.024* (0.014)		-0.093*** (0.023)	-0.046** (0.021)	-0.024 (0.020)
Skill Index		0.266*** (0.012)	0.067*** (0.009)			0.010*** (0.000)	0.005*** (0.000)
Log Previous Wage			0.598*** (0.016)				0.530*** (0.014)

Notes: Estimates are from regressions of the log of the real starting wage for the worker's current job. Results in the first set of columns correspond to estimates based on data simulated from our model. Results in the second set of columns correspond to estimates based on the SCE data, as reported in Table 2. For the simulated data regressions, we draw a sample size equal to that used in the real data. We repeat this 100 times and report the means and standard deviations of those estimates. Standard errors are presented in parentheses below the point estimates. *** denotes significance at the 1% level, ** denotes significance at the 5% level, and * denotes significance at the 10% level.

The results for job tenure are reported in Table 14. Overall the regressions in simulated and real data generate very similar results, both between tenure and the two types of referrals, and between tenure and previous wages.

Finally we look at how contact rates are related to the usage of referrals. In both the real and simulated data we find a strong positive relationship between contact rates and business referrals, with similar magnitudes. For family/friends referrals we find a negative relationship with contract rates in both cases, although the magnitudes in the simulated data are slightly smaller.

Table 14: Simulated Regressions Tenure

	<i>Model</i>				<i>Data</i>		
	Log Job Duration				Log Job Duration		
	(1)	(2)	(3)		(1)	(2)	(3)
Business Referral	-0.170*** (0.057)	-0.170*** (0.057)	-0.143** (0.056)		-0.185*** (0.055)	-0.195*** (0.054)	-0.172*** (0.065)
Family/Friends Referral	0.226*** (0.045)	0.227*** (0.046)	0.218*** (0.046)		0.216*** (0.049)	0.246*** (0.049)	0.229*** (0.060)
Skill Index		0.001 (0.039)	0.074* (0.040)			0.007*** (0.001)	0.007*** (0.001)
Log Previous Wage			-0.219*** (0.050)				-0.142*** (0.036)

Notes: Estimates are from regressions of the log of the duration of the current job. Results in the first set of columns correspond to estimates based on data simulated from our model. Results in the second set of columns correspond to estimates based on the SCE data, as reported in Table 3. For the simulated data regressions, we draw a sample size equal to that used in the real data. We repeat this 100 times and report the means and standard deviations of those estimates. Standard errors are presented in parentheses below the point estimates. *** denotes significance at the 1% level, ** denotes significance at the 5% level, and * denotes significance at the 10% level.

Table 15: Simulated Regressions Contact Rates

	<i>Model</i>				<i>Data</i>		
	Probability of Contact (Last 4 Weeks)				Probability of Contact (Last 4 Weeks)		
	(1)	(2)	(3)		(1)	(2)	(3)
Business Referral	0.064*** (0.020)	0.078*** (0.019)	0.033* (0.019)		0.050*** (0.018)	0.048*** (0.018)	0.037* (0.023)
Family/Friends Referral	-0.011 (0.017)	-0.031* (0.017)	-0.016 (0.017)		-0.043*** (0.016)	-0.037** (0.016)	-0.050** (0.021)
Skill Index		-0.235*** (0.013)	-0.362*** (0.014)			0.001*** (0.000)	0.001 (0.000)
Log Previous Wage			0.382*** (0.016)				0.047*** (0.013)

Notes: Estimates are from regressions of an indicator for whether or not an individual had contact with at least one potential employer in the last four weeks. Results in the first set of columns correspond to estimates based on data simulated from our model. Results in the second set of columns correspond to estimates based on the SCE data, as reported in Table 4. For the simulated data regressions, we draw a sample size equal to that used in the real data. We repeat this 100 times and report the means and standard deviations of those estimates. Standard errors are presented in parentheses below the point estimates. *** denotes significance at the 1% level, ** denotes significance at the 5% level, and * denotes significance at the 10% level.

References

- ALON, S. AND H. STIER (2019): “Job Search, Gender, and the Quality of Employment in Israel,” in *Stratification in Israel*, Routledge, 331–348.
- ARBEX, M., D. O’DEA, AND D. WICZER (2019): “Network Search: Climbing the Job Ladder Faster,” *International Economic Review*, 60, 693–720.
- BANDIERA, O., I. BARANKAY, AND I. RASUL (2009): “Social Connections and Incentives in the Workplace: Evidence from Personnel Data,” *Econometrica*, 77, 1047–1094.
- BAYER, P., S. L. ROSS, AND G. TOPA (2008): “Place of Work and Place of Residence: Informal Hiring Networks and Labor Market Outcomes,” *Journal of Political Economy*, 116, 1150–1196.
- BEAMAN, L. AND J. MAGRUDER (2012): “Who Gets the Job Referral? Evidence from a Social Networks Experiment,” *American Economic Review*, 102, 3574–3593.
- BENTOLILA, S., C. MICHELACCI, AND J. SUAREZ (2010): “Social Contacts and Occupational Choice,” *Economica*, 77, 20–45.
- BILAL, A., N. ENGBOM, S. MONGEY, AND G. L. VIOLANTE (2022): “Firm and Worker Dynamics in a Frictional Labor Market,” *Econometrica*, 90, 1425–1462.
- BOLTE, L., N. IMMORLICA, AND M. O. JACKSON (2020): “The Role of Referrals in Immobility, Inequality, and Inefficiency in Labor Markets,” *arXiv preprint arXiv:2012.15753*.
- BRIDGES, W. P. AND W. J. VILLEMEZ (1986): “Informal Hiring and Income in the Labor Market,” *American Sociological Review*, 51, 574–582.
- BROWN, M., E. SETREN, AND G. TOPA (2016): “Do Informal Referrals Lead to Better Matches? Evidence from a Firm’s Employee Referral System,” *Journal of Labor Economics*, 34, 161–209.
- BURKS, S. V., B. COWGILL, M. HOFFMAN, AND M. HOUSMAN (2015): “The Value of Hiring through Employee Referrals,” *Quarterly Journal of Economics*, 130, 805–839.
- CAHUC, P., F. POSTEL-VINAY, AND J.-M. ROBIN (2006): “Wage Bargaining with On-the-Job Search: Theory and Evidence,” *Econometrica*, 74, 323–364.
- CAIRÓ, I. AND T. CAJNER (2018): “Human Capital and Unemployment Dynamics: Why More Educated Workers Enjoy Greater Employment Stability,” *Economic Journal*, 128, 652–682.

- CALDWELL, S. AND N. HARMON (2019): “Outside Options, Bargaining, and Wages: Evidence from Coworker Networks,” *Working Paper*.
- CALVÓ-ARMENGOL, A. AND M. O. JACKSON (2004): “The Effects of Social Networks on Employment and Inequality,” *American Economic Review*, 94, 426–454.
- CASELLA, A. AND N. HANAKI (2008): “Information Channels in Labor Markets: On the Resilience of Referral Hiring,” *Journal of Economic Behavior & Organization*, 66, 492–513.
- CASTILLA, E. J. (2005): “Social Networks and Employee Performance in a Call Center,” *American Journal of Sociology*, 110, 1243–1283.
- CHETTY, R., N. HENDREN, AND L. F. KATZ (2016): “The Effects of Exposure to Better Neighborhoods on Children: New Evidence from the Moving to Opportunity Experiment,” *American Economic Review*, 106, 855–902.
- CORCORAN, M., L. DATCHER, AND G. DUNCAN (1980): “Information and Influence Networks in Labor Markets,” *Five Thousand American Families: Patterns of Economic Progress*, 8, 37.
- DATCHER, L. (1983): “The Impact of Informal Networks on Quit Behavior,” *Review of Economics and Statistics*, 65, 491–495.
- DIAMOND, R. (2016): “The Determinants and Welfare Implications of US Workers’ Diverging Location Choices by Skill: 1980-2000,” *American Economic Review*, 106, 479–524.
- DUSTMANN, C., A. GLITZ, U. SCHÖNBERG, AND H. BRÜCKER (2016): “Referral-based Job Search Networks,” *Review of Economic Studies*, 83, 514–546.
- ELLIOTT, J. R. (1999): “Social Isolation and Labor Market Insulation: Network and Neighborhood Effects on Less-Educated Urban Workers,” *The Sociological Quarterly*, 40, 199–216.
- ELSBY, M., A. GOTTFRIES, R. MICHAELS, AND D. RATNER (2022): “Vacancy Chains,” *Federal Reserve Bank of Philadelphia Working Paper Series*.
- ELSBY, M. W. AND A. GOTTFRIES (2022): “Firm Dynamics, On-the-Job Search, and Labor Market Fluctuations,” *Review of Economic Studies*, 89, 1370–1419.
- FABERMAN, R. J., A. I. MUELLER, A. ŞAHİN, AND G. TOPA (2022): “Job Search Behavior Among the Employed and Non-Employed,” *Econometrica*, 90, 1743–1779.
- FAFCHAMPS, M. AND A. MORADI (2015): “Referral and Job Performance: Evidence from the Ghana Colonial Army,” *Economic Development and Cultural Change*, 63, 715–751.

- FRIEBEL, G., M. HEINZ, M. HOFFMAN, AND N. ZUBANOV (2023): “What Do Employee Referral Programs Do? Measuring the Direct and Overall Effects of a Management Practice,” *Journal of Political Economy*, 131, 633–686.
- GALENIANOS, M. (2013): “Learning About Match Quality and the Use of Referrals,” *Review of Economic Dynamics*, 16, 668–690.
- (2014): “Hiring Through Referrals,” *Journal of Economic Theory*, 152, 304–323.
- GALEOTTI, A. AND L. P. MERLINO (2014): “Endogenous Job Contact Networks,” *International Economic Review*, 55, 1201–1226.
- GEE, L. K., J. JONES, AND M. BURKE (2017): “Social Networks and Labor Markets: How Strong Ties Relate to Job Finding on Facebook’s Social Network,” *Journal of Labor Economics*, 35, 485–518.
- GRANOVETTER, M. (1995): *Getting a Job: A Study of Contacts and Careers*, Chicago: University of Chicago Press.
- GREEN, P. G., L. M. TIGGES, AND I. BROWNE (1995): “Social Resources, Job Search, and Poverty in Atlanta,” *Research in Community Sociology*, 5, 161–182.
- GREGG, P. AND J. WADSWORTH (1996): “How Effective are State Employment Agencies? Jobcentre Use and Job Matching in Britain,” *Oxford Bulletin of Economics and Statistics*, 58, 443–467.
- GREGORY, V., G. MENZIO, AND D. G. WICZER (2021): “The Alpha Beta Gamma of the Labor Market,” *National Bureau of Economic Research Working Paper Series*.
- HAGEDORN, M. AND I. MANOVSKII (2008): “The Cyclical Behavior of Equilibrium Unemployment and Vacancies Revisited,” *American Economic Review*, 98, 1692–1706.
- HEATH, R. (2018): “Why Do Firms Hire Using Referrals? Evidence from Bangladeshi Garment Factories,” *Journal of Political Economy*, 126, 1691–1746.
- HECKMAN, J. J. AND B. SINGER (1984): “Econometric Duration Analysis,” *Journal of Econometrics*, 24, 63–132.
- HELLERSTEIN, J. K. AND D. NEUMARK (2020): “Social Capital, Networks, and Economic Wellbeing,” *IZA Discussion Papers*.
- HENSVIK, L. AND O. N. SKANS (2016): “Social Networks, Employee Selection, and Labor Market Outcomes,” *Journal of Labor Economics*, 34, 825–867.

- HOLZER, H. J. (1987a): "Hiring Procedures in the Firm: Their Economic Determinants and Outcomes," Tech. rep., National Bureau of Economic Research.
- (1987b): "Job Search by Employed and Unemployed Youth," *ILR Review*, 40, 601–611.
- (1988): "Search Method Use by Unemployed Youth," *Journal of Labor Economics*, 6, 1–20.
- IOANNIDES, Y. M. AND L. D. LOURY (2004): "Job Information Networks, Neighborhood Effects, and Inequality," *Journal of Economic Literature*, 42, 1056–1093.
- JOVANOVIĆ, B. (1979): "Job Matching and the Theory of Turnover," *Journal of Political Economy*, 87, 972–990.
- KATZ, L. F., J. R. KLING, AND J. B. LIEBMAN (2001): "Moving to Opportunity in Boston: Early Results of a Randomized Mobility Experiment," *Quarterly Journal of Economics*, 116, 607–654.
- KENNAN, J. AND J. R. WALKER (2011): "The Effect of Expected Income on Individual Migration Decisions," *Econometrica*, 79, 211–251.
- KORENMAN, S. AND S. C. TURNER (1996): "Employment Contacts and Minority-White Wage Differences," *Industrial Relations: A Journal of Economy and Society*, 35, 106–122.
- KRAMARZ, F. AND O. N. SKANS (2014): "When Strong Ties are Strong: Networks and Youth Labour Market Entry," *Review of Economic Studies*, 81, 1164–1200.
- KUGLER, A. D. (2003): "Employee Referrals and Efficiency Wages," *Labour Economics*, 10, 531–556.
- LIN, N., W. M. ENSEL, AND J. C. VAUGHN (1981): "Social Resources and Strength of Ties: Structural Factors in Occupational Status Attainment," *American Sociological Review*, 46, 393–405.
- LISE, J., C. MEGHIR, AND J.-M. ROBIN (2016): "Matching, Sorting and Wages," *Review of Economic Dynamics*, 19, 63–87.
- LOURY, L. D. (2006): "Some Contacts Are More Equal than Others: Informal Networks, Job Tenure, and Wages," *Journal of Labor Economics*, 24, 299–318.
- MARMAROS, D. AND B. SACERDOTE (2002): "Peer and Social Networks in Job Search," *European Economic Review*, 46, 870–879.

- MARSDEN, P. V. (2017): “Interpersonal Ties, Social Capital, and Employer Staffing Practices,” in *Social Capital*, Routledge, 105–125.
- MARSDEN, P. V. AND E. H. GORMAN (2001): “Social Networks, Job Changes, and Recruitment,” *Sourcebook of Labor Markets*, 467–502.
- MILLER, S. R. AND J. E. ROSENBAUM (1997): “Hiring in a Hobbesian World: Social Infrastructure and Employers’ Use of Information,” *Work and Occupations*, 24, 498–523.
- MONTGOMERY, J. D. (1991): “Social Networks and Labor-Market Outcomes: Toward an Economic Analysis,” *American Economic Review*, 81, 1408–1418.
- MOON, J.-W. (2023): “Strategic Referrals and On-the-Job Search Equilibrium,” *Journal of Monetary Economics*, 134, 135–151.
- MORTENSEN, D. T. AND C. A. PISSARIDES (1994): “Job Creation and Job Destruction in the Theory of Unemployment,” *The Review of Economic Studies*, 61, 397–415.
- MOSCARINI, G. AND F. POSTEL-VINAY (2023): “On the Job Search and Business Cycles,” *Revue Economique*, Forthcoming.
- MOUW, T. (2003): “Social Capital and Finding a Job: Do Contacts Matter?” *American Sociological Review*, 68, 868–898.
- MUKOYAMA, T., C. PATTERSON, AND A. ŞAHİN (2018): “Job Search Behavior over the Business Cycle,” *American Economic Journal: Macroeconomics*, 10, 190–215.
- NECKERMAN, K. M. AND J. KIRSCHENMAN (1991): “Hiring Strategies, Racial Bias, and Inner-City Workers,” *Social Problems*, 38, 433–447.
- PALLAIS, A. AND E. G. SANDS (2016): “Why the Referential Treatment? Evidence from Field Experiments on Referrals,” *Journal of Political Economy*, 124, 1793–1828.
- PELLIZZARI, M. (2010): “Do Friends and Relatives Really Help in Getting a Good Job?” *ILR Review*, 63, 494–510.
- PISTAFERRI, L. (1999): “Informal Networks in the Italian Labor Market,” *Giornale degli Economisti e Annali di Economia*, 355–375.
- PIYAPROMDEE, S. (2021): “The Impact of Immigration on Wages, Internal Migration, and Welfare,” *Review of Economic Studies*, 88, 406–453.
- POSTEL-VINAY, F. AND J.-M. ROBIN (2002): “Equilibrium Wage Dispersion with Worker and Employer Heterogeneity,” *Econometrica*, 70, 2295–2350.

- SAN, S. (2021): “Who Works Where and Why? Parental Networks and the Labor Market,” *Working Paper*.
- SCHMUTTE, I. M. (2015): “Job Referral Networks and the Determination of Earnings in Local Labor Markets,” *Journal of Labor Economics*, 33, 1–32.
- SHIMER, R. (2005): “The Cyclical Behavior of Equilibrium Unemployment and Vacancies,” *American Economic Review*, 95, 25–49.
- SIMON, C. J. AND J. T. WARNER (1992): “Matchmaker, Matchmaker: The Effect of Old Boy Networks on Job Match Quality, Earnings, and Tenure,” *Journal of Labor Economics*, 10, 306–330.
- TOPA, G. (2001): “Social Interactions, Local Spillovers and Unemployment,” *Review of Economic Studies*, 68, 261–295.
- (2011): “Labor Markets and Referrals,” in *Handbook of Social Economics*, Elsevier, vol. 1, 1193–1221.
- (2019): “Social and Spatial Networks in Labour Markets,” *Oxford Review of Economic Policy*, 35, 722–745.
- WAHBA, J. AND Y. ZENOU (2005): “Density, Social Networks and Job Search Methods: Theory and Application to Egypt,” *Journal of Development Economics*, 78, 443–473.
- ZERECERO, M. (2021): “The Birthplace Premium,” *Working Paper*.